

Final Report


NASA Research Grant No. NAG-1-968

(NASA-CR-190529) SPACECRAFT
DYNAMICS CHARACTERIZATION AND
CONTROL SYSTEM FAILURE DETECTION.
VOLUME 3: CONTROL SYSTEM FAILURE
MONITORING Final Report (MIT)
140 p


N92-31254

Unclass

63/63 0109012



**DEPARTMENT OF
AERONAUTICS AND ASTRONAUTICS
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
CAMBRIDGE, MASSACHUSETTS 02139**



SPACECRAFT DYNAMICS CHARACTERIZATION AND CONTROL SYSTEM FAILURE DETECTION

LANGLEY REPORT

IN-63-OR

Final Report

109012

NASA Research Grant No. NAG-1-968

p.140

Volume III

CONTROL SYSTEM FAILURE MONITORING

Christiaan M. van Schalkwyk

*Department of Aeronautics and Astronautics
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139*

February 10, 1992

Contents

List of Figures	5
Notation	9
Symbols	10
1 Introduction	12
2 Generalized Parity Relations	16
2.1 Single Sensor Parity Relations	18
2.2 Double Sensor Parity Relations	23
2.3 Actuator Parity Relations	26
2.4 Example	29
3 Displacement Sensor Failure Detection	35
3.1 Introduction	35
3.2 Model-based Single Sensor Parity Relations	41

3.3	Identified Single Sensor Parity Relations	47
3.4	Transfer functions of model-based and identified Single Sensor Parity Relations	51
3.5	Increased Sampling Period	55
3.6	Double Sensor Parity Relations	57
3.7	Summary	68
4	Accelerometer and Gyro Failure Detection	70
4.1	Introduction	70
4.2	Single Sensor Parity Relations	74
4.3	Double Sensor Parity Relations	79
4.4	Summary	87
5	Actuator Failure Detection	88
6	Generalized Parity Relation Experiments Using the ACES Facility	96
6.1	Introduction	96
6.2	Processing the residuals	98
6.3	ACES Failure Experiments	99
6.4	Summary	108

7	Failure Detection Filters	109
7.1	Introduction	109
7.2	Actuator Failure Experiments	114
7.3	Sensor Failure Experiments	121
7.4	Computational Burden	123
7.5	Summary	126
8	Conclusion	128
A	Second order system analysis	133
B	Mini-Mast state-space model	135
	Bibliography	138

List of Figures

1.1	FDI block diagram	13
2.1	Block diagram of SSPR Residual Generator	23
2.2	Block diagram of DSPR Residual Generator	25
2.3	Block diagram of SAPR Residual Generator	29
2.4	Block diagram of the plant and SSPR Residual Generator	32
2.5	Transfer functions of the SSPR Residual Generator	33
3.1	Schematic diagram of the Mini-Mast.	36
3.2	Spectrum of input signal	39
3.3	Spectrum of Displacement Sensor 2	40
3.4	Displacement Sensor D1 failure	42
3.5	Components $r_{1,y}$ and $r_{1,u}$ of SSPR r_1	43
3.6	Displacement Sensor D3 failures	45
3.7	Noisy sensor	46
3.8	Identified SSPR residual r_3	47
3.9	Components of identified SSPR residual r_3	49
3.10	Identified SSPR r_3 , 20 lags	50

3.11 SSPR Residual Generator zero locations	52
3.12 SSPR Residual Generator transfer functions	53
3.13 Spectra of SSPR residuals	54
3.14 Identified SSPR r_3 , 10 lags, $2T_s$	55
3.15 Identified SSPR r_3 , 20 lags, $2T_s$	56
3.16 Identified SSPR Residual Generator transfer functions $1T_s$ and $2T_s$	57
3.17 Model-based DSPR residuals r_{12} and r_{13}	59
3.18 Model-based DSPR residuals r_{12} and r_{13}	60
3.19 Identified DSPR residuals r_{12} and r_{32}	61
3.20 Dual identified DSPR residuals r_{21} and r_{23}	62
3.21 Identified DSPR residuals r_{13} and r_{23}	64
3.22 Identified DSPR residuals r_{21} and r_{23} , 20 lags	65
3.23 Model-based DSPR residuals r_{12} and r_{13} , 10 lags, $2T_s$	66
3.24 Identified DSPR residuals r_{21} and r_{23} , 20 lags, $2T_s$	67
4.1 Spectrum of the Y-axis acceleration, u = filtered white noise	71
4.2 Spectrum of the Y-axis acceleration, u = $4T_s$ white noise	72
4.3 Spectrum of the Z-axis gyro signal, u = $4T_s$ white noise	73
4.4 Experimental setup	74
4.5 SSPR for Y-axis accelerometer failure, 20 lags, $1T_s$	75
4.6 SSPR for Z-axis gyro failure, 20 lags, $1T_s$	75
4.7 SSPR for Y-axis accelerometer failure, 20 lags, $1T_s$; filtered residual	76
4.8 SSPR for Z-axis gyro failure, 20 lags, $1T_s$; filtered residual	76

4.9	SSPR for Y -axis accelerometer failure, 20 lags, $2T_s$	77
4.10	SSPR for Z -axis gyro failure, 20 lags, $2T_s$	77
4.11	DSPR for X -axis accelerometer failure, (11,10) lags, $1T_s$	80
4.12	DSPR for Z -axis gyro failure, (11,10) lags, $1T_s$	81
4.13	DSPR for X -axis accelerometer failure, (11,10) lags, $2T_s$	82
4.14	DSPR for Z -axis gyro failure, (11,10) lags, $2T_s$	83
4.15	DSPR for X -axis accelerometer failure, (11,10) lags, $2T_s$	84
4.16	SSPR for Z -axis gyro failure, 20 lags, $1T_s$	86
4.17	SSPR for Z -axis gyro failure, 20 lags, $1T_s$	86
5.1	Model-based SAPR for X -wheel failure	89
5.2	Identified SAPR for X -wheel failure	89
5.3	SAPR residual for Y -torque wheel failure	91
5.4	SAPR residual for Y -torque wheel failure, noisy measurement	92
5.5	SAPR residual contributions for Y -torque wheel failure, noisy measurement	93
5.6	SAPR Residual Generator transfer functions	94
6.1	Schematic diagram of the ACES mast	97
6.2	Block diagram of post processing filter.	98
6.3	SSPR residual for base gyro	101
6.4	SSPR residual for LMED accelerometer	101
6.5	DSPR residual for (base gyro, tip gyro) pair	103
6.6	DSPR residual for (base gyro, tip gyro) pair	103

6.7	DSPR residual for (base gyro, tip accelerometer) pair	104
6.8	DSPR residual for (base gyro, tip accelerometer) pair	104
6.9	DSPR residual for (tip gyro, LMED accelerometer) pair	105
6.10	DSPR residual for (base gyro, tip accelerometer) pair	106
6.11	SAPR individual contributions for X -gimbal torquer	107
7.1	Block diagram of post processing filter.	115
7.2	FDF residual for X -torque wheel failure: r_{f1}	116
7.3	FDF residual for X -torque wheel failure: r_{p1}	116
7.4	FDF residual for Y -torque wheel failure: r_{p2}	118
7.5	FDF residual for Z -torque wheel failure	118
7.6	SAPR residual for Y -torque wheel failure	119
7.7	SAPR transfer functions	120
7.8	Single actuator FDF transfer functions	120
7.9	FDF residual for Sensor D1 failure	121
7.10	FDF residual for Sensor D1 failure	122
7.11	FDF residual for Y -torque wheel failure	124
7.12	Unprocessed FDF residual for Y -torque wheel failure	124

Notation

Vector notation

Vectors will be written as a single column and will be denoted by bold lowercase characters. We will also use the Matlab notation where a column vector is written on one line and a semicolon is used to delimit the elements of the vector:

$$\mathbf{x} \in \mathbb{R}^n \iff \mathbf{x} = \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} \quad (0.1)$$

$$= [x_1; \dots; x_n]. \quad (0.2)$$

Row vectors will be written on one line and a comma will be used to delimit the elements of the vector:

$$\mathbf{x}' \in \mathbb{R}^{1 \times n} \iff \mathbf{x}' = [x_1, \dots, x_n]. \quad (0.3)$$

A few special vectors, which will be defined explicitly, will be written with reversed indices:

$$\boldsymbol{\beta} \in \mathbb{R}^{l+1} \iff \boldsymbol{\beta} = [\beta_l; \dots; \beta_0]. \quad (0.4)$$

Matrix notation

Bold uppercase letters will be used to denote matrices, the corresponding lowercase letters with subscripts ij will be used to denote the (i, j) entry:

$$\mathbf{A} \in \mathbb{R}^{m \times n} \iff \mathbf{A} = \begin{pmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & & \vdots \\ a_{m1} & \cdots & a_{mn} \end{pmatrix}, \quad a_{ij} \in \mathbb{R}. \quad (0.5)$$

Columns of the matrix will be denoted by the vectors $\mathbf{a}_1, \dots, \mathbf{a}_n$, and the rows will be denoted by the row vectors $\mathbf{a}'_1, \dots, \mathbf{a}'_m$. The transpose of the matrix will be written as \mathbf{A}^T .

Symbols

α_i	parity relation coefficients, Equations (2.23) and (2.54)
α_{ij}	parity relation coefficients, Equations (2.36)
A	discrete-time state transition matrix, Equation (2.8)
A_c	continuous-time system matrix, Equation (2.5)
A_f	detection filter system matrix, Equation (7.3)
A_p	system matrix, Equation (7.1)
A_1, A_2	system matrices, Equations (7.19) and (7.21)
β_i	parity relation coefficients, Equations (2.18) and (2.58)
B	discrete-time input matrix, Equation (2.8)
B_c	continuous-time input matrix, Equation (2.5)
B_f	detection filter input matrix, Equation (7.3)
B_{fu}, B_{fy}	detection filter input matrices, Equation (7.14)
B_p	input matrix, Equation (7.1)
B_1, B_2	input matrices, Equations (7.19) and (7.21)
C	output matrix, Equation (2.9)
C_f	detection filter output matrix, Equation (7.15)
C_p	output matrix, Equation (7.2)
C_1, C_2	output matrices, Equations (7.19) and (7.21)
c'_i	i th row of C
C_i	Equation (2.16)
C_{ij}	Equation (2.32)
\bar{C}_i	Equation (2.52)
D	feedforward matrix, Equation (2.9), detector gain, Equation (7.3)
D_{fy}	detection filter feedforward matrix, Equation (7.15)
d'_i	i th row of D
D_i	Equation (2.17)
D_{ij}	Equation (2.33)
\bar{D}_i	Equation (2.53)
$e(t)$	error signal, Equation (7.4)

J_1, J_2	Equation (7.18)
L_1, L_2	Equation (7.16)
l_1, l_2	Equation (7.12)
$m_1(t), m_2(t)$	arbitrary functions of time, Equation (7.12)
$\mathbf{m}_1(t), \mathbf{m}_2(t)$	arbitrary functions of time, Equation (7.16)
$n(t)$	arbitrary function of time, Equation (7.10)
$\mathbf{n}_1(t), \mathbf{n}_2(t)$	arbitrary functions of time, Equation (7.18)
n_i	Equations (2.12), (2.28), (2.40)
n_j	Equation (2.28)
\mathbb{R}	field of real numbers
$r_i(k)$	i th SSPR or SAPR residual, Equations (2.21), (2.56)
$r_{ij}(k)$	ij th DSPR or DAPR residual, Equation (2.37)
$r_{ji}(k)$	ji th DSPR or DAPR residual, Equation (2.38)
T_s	sampling period
$\mathbf{u}(t)$	continuous-time input vector, Equation (2.5)
$\mathbf{u}(k)$	discrete-time input vector, Equation (2.8)
$u_i(k)$	i th element of $\mathbf{u}(k)$
$\mathbf{u}_a(t)$	actual input vector, Equation (7.1)
$\mathbf{u}_c(t)$	commanded input vector, Equation (7.3)
$\mathbf{u}(n_i)$	Equations (2.15) and (2.41)
$\mathbf{u}_i(n_i)$	Equation (2.50)
$\mathbf{x}(t)$	continuous-time state vector, Equation (2.5)
$\mathbf{x}(k)$	discrete-time state vector, Equation (2.8)
$\mathbf{x}_p(t)$	state vector, Equation (7.1)
$\mathbf{x}_{n_1}, \mathbf{x}_{n_2}$	state vectors, Equations (7.19) and (7.21)
$\mathbf{y}(t)$	continuous-time measurement vector, Equation (2.6)
$\mathbf{y}(k)$	discrete-time measurement vector, Equation (2.9)
$y_i(k)$	i th element of $\mathbf{y}(k)$
$\mathbf{y}(n_i)$	Equation (2.40)
$\mathbf{y}_i(n_i)$	Equation (2.14)
$\mathbf{z}(t)$	detection filter state vector, Equation (7.3)

Chapter 1

Introduction

The requirement that a control system must be tolerant to the failure of its components and still perform safely and reliably puts stringent requirements on the reliability of the components that are used. Often the requirements on the reliability are so strict that it can only be achieved through some form of redundancy. An example is flexible space structures. Due to their large sizes and lightweight construction they have very low damping so that active control is necessary to do shape control and damp out vibrations throughout the structure. Active control is also necessary to perform other tasks like stationkeeping and attitude control. Systems in space must work for long unattended periods of time and with long intervals between maintenance so that a control system must be able to perform satisfactorily even when some of its components, especially the actuators and sensors, fail. To ensure stability of the control system and continue the mission it is necessary to detect the failure of a component. Once a failure has been detected and the failed component has been identified, the control system must be reconfigured to isolate the faulty component from the controller. Other examples of control systems that require very high reliability are

aircraft engines, nuclear reactors, and process control systems, to name but a few.

To increase the reliability of a system some form of redundancy is usually used. Redundancy can be divided into two classes, hardware redundancy and analytical redundancy. In hardware redundancy the reliability is increased by replicating the control system components. A solution that is often applied is to use three or more sensors of the same kind to measure the same variable. A voting scheme is then employed to find the odd one out. Hardware redundancy has the advantage that it is insensitive to the magnitude of the failure and can detect any type of discrepancy. Although hardware redundancy is simple to implement, it is costly and adds unnecessary weight to the system. When many sensors and actuators are used it becomes impractical to triplicate each device. As an example, it is estimated that a large flexible space structure will have approximately 200 control system components. Tripling so many components is impractical and not cost effective. Another way to increase the reliability of a system is through *analytical redundancy*. Here the redundancy present in the model of the plant and input-output histories are used to detect and identify the failure of a component.

The typical form of a failure detection and isolation (FDI) system is shown in Figure 1.1. The FDI system is divided into two subsystems, the generation of residuals and decision




Figure 1.1: FDI block diagram.

making, as shown in the figure. The Residual Generator uses the commanded inputs to the plant, the measured outputs from the plant, and a model of the plant to generate a set of residuals. The generation of residuals has been studied for many years and surveys of these methods can be found in Willsky [14], Basseville [1], and Merrill [12]. The Decision Function analyzes the residuals and based on this analysis makes a decision about the state of the actuators and sensors. Typical examples of this analysis are simple threshold detectors that compare the magnitudes of the residuals with a set of thresholds and declaring a failure when the amplitude exceeds the threshold. Other methods are moving average analysis and statistical decision theory. In the latter case *a priori* probabilities of the failure modes are hypothesized and it is possible to optimize for a specific mode of failure. It is not always possible to enumerate all modes of failure and obtain the corresponding probabilities. It is therefore desirable to have a methodology that does not require the specification of the failure modes and corresponding probabilities of failure. Also, the method should be applicable to both sensors and actuators. Only two methods satisfy the requirements set forth, the Failure Detection Filter by Beard [2] (see also Jones [6] and Massoumnia [10]) and the method of Generalized Parity Relations by Chow [4]. Because all analytical redundancy methods use a model of the plant they are all sensitive to modelling errors. The design of robust parity relations has been discussed by Lou et al. [8].

In this work we discuss the application of Generalized Parity Relations and the Failure Detection Filter to two experimental flexible space structures, the NASA Langley Mini-Mast and Marshall Space Flight Center ACES mast. We concentrated on the generation of residuals and made no attempt to implement the Decision Function. It should be clear from the examples that are presented in later chapters whether it would be possible to detect the failure of a specific component. The report is structured as follows. In Chapter 2 we derive the equations for Generalized Parity Relations. Two special cases are treated: namely,

Single Sensor Parity Relations (SSPR) and Double Sensor Parity Relations (DSPR). Generalized Parity Relations for actuators are also derived. Chapter 3 describes the NASA Langley Mini-Mast and discusses the application of SSPR and DSPR to a set of displacement sensors located at the tip of the Mini-Mast. The performance of a reduced order model that includes the first five modes of the mast is compared to a set of parity relations that was identified on a set of input-output data. Both time domain and frequency domain comparisons are made. The effect of the sampling period and model order on the performance of the Residual Generators are also discussed. Chapter 4 presents failure detection experiments where the sensor set consisted of two gyros and an accelerometer. The effects of model order and sampling frequency are again illustrated. The detection of actuator failures are discussed in Chapter 5. In Chapter 6 we use Generalized Parity Relations to monitor control system component failures on the ACES mast. Chapter 7 gives an overview of the Failure Detection Filter and experimental results are then discussed. Conclusions and directions for future research are given in Chapter 8.

Chapter 2

Generalized Parity Relations

In the previous chapter we gave an outline of an FDI system where, for convenience of analysis, we divided the system into two functional parts: the Residual Generator and the Decision Function. In this chapter we give a brief description of a method to generate residuals. The method, known as Generalized Parity Relations, is treated in detail by Chow [4] and Dutilloy [5].

There are two forms of analytical redundancy, namely direct redundancy and temporal redundancy. In direct redundancy a relation is formed by taking a linear combination of the instantaneous values of a set of sensors whose outputs are linearly dependent. As an example, let I denote a set of sensors whose instantaneous outputs are linearly dependent and let the j th sensor be a member of the set. We can then find a relation for the j th output y_j :

$$y_j(t) = \sum_{\substack{i \in I \\ i \neq j}} \alpha_i y_i(t). \quad (2.1)$$

The residual is then defined as

$$r(t) = y_j(t) - \sum_{\substack{i \in I \\ i \neq j}} \alpha_i y_i(t) \quad (2.2)$$

which will be zero (except for noise or other unmodelled effects) when all the sensors are fully operational and nonzero in the case of a failure. Note that if $r(t)$ is nonzero, any of the sensors in the set could have failed — this single relation does not indicate which sensor has failed.

In temporal redundancy, the histories of outputs and inputs are taken into account. The following example is used to illustrate temporal redundancy: consider a vehicle with mass m and velocity $v(t)$ with commanded force $f(t)$ being applied to it. The velocity at time $t + \Delta t$ is given by the relation

$$v(t + \Delta t) = v(t) + \frac{f(t)}{m} \Delta t. \quad (2.3)$$

The velocity measurements $v(t)$ and $v(t + \Delta t)$ are now used together with the commanded force to form the residual

$$r(t + \Delta t) = v(t + \Delta t) - v(t) - \frac{f(t)}{m} \Delta t. \quad (2.4)$$

If the rate sensor fails in some way the measured velocity will differ from the actual velocity so that residual $r(t + \Delta t)$ will be nonzero. Thus, the nonzero residual indicates the failure of the sensor. When the actuator fails, the force applied to the mass will be different from the commanded force that is used to compute the residual. Hence, the residual will be nonzero and we can also detect the failure of the actuator. In this example, both the sensor failure and the actuator failure result in the residual being nonzero; therefore, without additional information we cannot determine which one of the components has failed when we observe a nonzero residual.

In our discussion so far we assumed that the residual is exactly zero when the system is in perfect working condition — in a practical FDI system this will never be the case because there will always be measurement noise, disturbances, and model mismatches. For the example under discussion, the only parameter for the plant is the mass m and, for the residual to have a small amplitude, the mass must be known accurately. The best we can hope for in a practical system is a residual with a *small* amplitude when all the components are functional and a *large* amplitude when a component has failed. Hopefully the difference between *small* and *large* will be large enough so that a threshold detector can then be used to discriminate between the failed and unfailed states. This example illustrates that generalized parity relations can be used to detect sensor and actuator failures and that the residual generator depends on the fidelity of the model to give a small residual when all the components are fully operational.

In this work we will discuss only temporal redundancy relations. Furthermore, the formulation of parity relations does not require the specification of measurement and process noise models; therefore, we will not include noise in the plant model. Chow [4] treated the case where noise is present in the system and discussed methods to obtain robust relations.

2.1 Single Sensor Parity Relations

Generalized parity relations can be constructed so that it is possible to identify which sensor has failed. The procedure is to construct parity relations from different subsets of the sensors so that when a sensor fails, only a subset of the parity residuals becomes larger. In this section we will discuss a specific method that can detect and identify sensor failures. The method, known as single sensor parity relations (SSPR), is discussed in detail by Dutilloy [5] and Massoumnia and Vander Velde [11]. The basic idea is to construct a set of relations

$\{r_i, i = 1, 2, \dots\}$ so that each residual r_i depends on one and only one sensor y_i . When a sensor fails only the corresponding residual is affected, and it is therefore very easy to identify which sensor has failed. In general, when an actuator fails, all the single sensor parity relations will be affected. In this case, the Decision Function (see Chapter 1) will decide that it was not all the sensors that have failed simultaneously as this is unlikely to happen.

We will assume that the plant can be modelled accurately by a continuous-time, linear, time-invariant model given by

$$\dot{\mathbf{x}}(t) = \mathbf{A}_c \mathbf{x}(t) + \mathbf{B}_c \mathbf{u}(t), \quad (2.5)$$

$$\mathbf{y}(t) = \mathbf{C} \mathbf{x}(t) + \mathbf{D} \mathbf{u}(t), \quad (2.6)$$

where $\mathbf{x}(t) \in \mathbb{R}^{n_x}$ is the state vector, $\mathbf{u}(t) \in \mathbb{R}^{n_u}$ is the commanded input vector, $\mathbf{y}(t) \in \mathbb{R}^{n_y}$ is the measurement vector, and $\mathbf{A}_c \in \mathbb{R}^{n_x \times n_x}$, $\mathbf{B}_c \in \mathbb{R}^{n_x \times n_u}$, $\mathbf{C} \in \mathbb{R}^{n_y \times n_x}$, and $\mathbf{D} \in \mathbb{R}^{n_y \times n_u}$ are the usual continuous-time state-space matrices. When a sensor fails the output can be modelled by

$$\mathbf{y}(t) = \mathbf{C} \mathbf{x}(t) + \mathbf{D} \mathbf{u}(t) + \mathbf{f}(t), \quad (2.7)$$

where the vector $\mathbf{f}(t)$ is an unknown function of time. This simple model is adequate to describe many failures that occur in practical systems and is discussed in more detail by Jones [6] and Massoumnia [10]. We will make no attempt to characterize $\mathbf{f}(t)$; an important property of generalized parity relations is that no failure modes and corresponding probabilities of failure need to be specified. It is important to notice that the output given by Equation (2.6) is modified in some sense when a sensor fails.

The construction of generalized parity relations requires a discrete-time model of the system. Let T_s denote the sampling period. If the input signal $\mathbf{u}(t)$ is constant over the

interval $kT_s \leq t < (k+1)T_s$, the continuous-time system of Equations (2.5) and (2.6) can be discretized as follows:

$$\begin{aligned} \mathbf{x}((k+1)T_s) &= e^{A_c T_s} \mathbf{x}(k) + \int_0^{T_s} e^{A_c(T_s-\tau)} \mathbf{B}_c d\tau \mathbf{u}(kT_s) \\ &= \mathbf{A} \mathbf{x}(k) + \mathbf{B} \mathbf{u}(kT_s), \end{aligned} \quad (2.8)$$

$$\mathbf{y}(kT_s) = \mathbf{C} \mathbf{x}(kT_s) + \mathbf{D} \mathbf{u}(kT_s), \quad (2.9)$$

where

$$\mathbf{A} = e^{A_c T_s}, \quad (2.10)$$

$$\mathbf{B} = \int_0^{T_s} e^{A_c(T_s-\tau)} \mathbf{B}_c d\tau \quad (2.11)$$

The notation $\mathbf{x}(k)$, $\mathbf{y}(k)$ and $\mathbf{u}(k)$ will often be used to denote $\mathbf{x}(kT_s)$, $\mathbf{y}(kT_s)$ and $\mathbf{u}(kT_s)$ respectively.

Consider now the i th sensor output y_i and let \mathbf{c}'_i and \mathbf{d}'_i denote the i th row of \mathbf{C} and \mathbf{D} respectively; the output history is easily obtained in terms of the initial state $\mathbf{x}(k)$ and inputs $\mathbf{u}(k)$, $\mathbf{u}(k+1)$, ... as

$$\begin{aligned} y_i(k) &= \mathbf{c}'_i \mathbf{x}(k) + \mathbf{d}'_i \mathbf{u}(k), \\ y_i(k+1) &= \mathbf{c}'_i \mathbf{A} \mathbf{x}(k) + \mathbf{c}'_i \mathbf{B} \mathbf{u}(k) + \mathbf{d}'_i \mathbf{u}(k+1), \\ y_i(k+2) &= \mathbf{c}'_i \mathbf{A}^2 \mathbf{x}(k) + \mathbf{c}'_i \mathbf{A} \mathbf{B} \mathbf{u}(k) + \mathbf{c}'_i \mathbf{B} \mathbf{u}(k+1) + \mathbf{d}'_i \mathbf{u}(k+2), \\ &\vdots \\ y_i(k+n_i) &= \mathbf{c}'_i \mathbf{A}^{n_i} \mathbf{x}(k) + \mathbf{c}'_i \mathbf{A}^{n_i-1} \mathbf{B} \mathbf{u}(k) + \cdots + \mathbf{c}'_i \mathbf{B} \mathbf{u}(k+n_i-1) + \mathbf{d}'_i \mathbf{u}(k+n_i). \end{aligned} \quad (2.12)$$

These equations can be written in a compact form as follows:

$$y_i(n_i) = C_i \mathbf{x}(k) + D_i \mathbf{u}(n_i), \quad (2.13)$$

where

$$y_i(n_i) = [y_i(k); y_i(k+1); \dots; y_i(k+n_i)], \quad (2.14)$$

$$u(n_i) = [u(k); u(k+1); \dots; u(k+n_i)], \quad (2.15)$$

$$C_i = [c'_i; c'_i A; \dots; c'_i A^{n_i}], \quad (2.16)$$

$$D_i = \begin{pmatrix} d'_i & 0 & 0 & \dots & 0 \\ c'_i B & d'_i & 0 & \dots & 0 \\ c'_i AB & c'_i B & d'_i & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ c'_i A^{n_i-1} B & c'_i A^{n_i-2} B & c'_i A^{n_i-3} B & \dots & d'_i \end{pmatrix}, \quad (2.17)$$

with $y_i \in \mathbb{R}^{n_i+1}$, $u \in \mathbb{R}^{(n_i+1)n_u}$, $C_i \in \mathbb{R}^{(n_i+1) \times n_x}$ and $D_i \in \mathbb{R}^{(n_i+1) \times (n_i+1)n_u}$. Note that the Cayley-Hamilton theorem assures that C_i will be singular for $n_i \geq n_x$. If n_i is chosen large enough so that the matrix C_i becomes singular, we can find a vector $\beta_i \in \mathbb{R}^{n_i+1}$ in the left null space of C_i so that

$$\beta_i^T C_i = 0, \quad (2.18)$$

$$\beta_i = [\beta_{i,n_i}; \beta_{i,n_i-1}; \dots; \beta_{i,1}; 1], \quad (2.19)$$

where we have scaled the vector so that last element, $\beta_0 = 1$. The reason for this choice will become clear later. If the system is observable from the i th sensor, $n_i = n_x$.

Multiplying Equation (2.13) by β_i^T and rearranging we get

$$\beta_i^T y_i(n_i) - \beta_i^T D_i u(n_i) = 0. \quad (2.20)$$

Equation (2.20) is called the i th single sensor parity relation. When the i th sensor fails, the output equation is modified in some unknown way so that the above relation will not hold.

We define the i th residual as

$$\begin{aligned} r_i(k + n_i) &= \beta_i^T y_i(n_i) - \beta_i^T D_i u(n_i) \\ &= \beta_i^T y_i(n_i) - \alpha_i^T u(n_i) \end{aligned} \quad (2.21)$$

$$= r_{i,y} - r_{i,u} \quad (2.22)$$

where $r_{i,y}$ is the contribution of the i th output, $r_{i,u}$ is the contribution of all the inputs and

$$\alpha_i = \beta_i^T D_i \quad (2.23)$$

$$\begin{aligned} &= [\alpha_{i,1,n_i}; \alpha_{i,2,n_i}; \dots; \alpha_{i,n_u,n_i}; \alpha_{i,1,n_i-1}; \alpha_{i,2,n_i-1}; \dots; \alpha_{i,n_u,n_i-1}; \dots; \\ &\quad \alpha_{i,1,0}; \alpha_{i,2,0}; \dots; \alpha_{i,n_u,0}], \end{aligned} \quad (2.24)$$

$\alpha_i \in \mathbb{R}^{(n_i+1)n_u}$. When all the sensors and actuators are fully operational, the model matches the plant exactly, and there are no measurement noise and disturbances, all the residuals r_i , $i = 1, 2, \dots, n_y$ will be zero. When the i th sensor fails, $r_i(k)$ will be nonzero and because the residuals $r_j(k)$, $j \neq i$, are not functions of the i th sensor, they will remain zero. Thus it is possible to detect and identify the failure of the i th sensor. Equation (2.21) has the form of a multi-input single-output finite impulse response filter and both the system input vector $u(k)$ and the scalar output $y_i(k)$ are inputs to the residual generator. A block diagram of the SSPR Residual Generator is shown in Figure 2.1. Because the system under discussion is time-invariant the starting time is arbitrary. Using this property and Equations (2.19) and (2.24), we can rewrite Equation (2.20) as summations,

$$\sum_{s=0}^{n_i} \beta_{i,s} y_i(k-s) = \sum_{r=1}^{n_u} \sum_{s=0}^{n_i} \alpha_{i,r,s} u_r(k-s) \quad (2.25)$$

which is an ARX model for the system. (ARX = autoregressive with external input.) The ARX description motivated the choice for $\beta_0 = 1$ as this gives a monic denominator polynomial for a single-input single-output system. If we can find an ARX model for the




Figure 2.1: Block diagram of SSPR Residual Generator.

plant we do not need to find the state-space matrices. Many system identification techniques immediately identify an ARX model from input-output data; see for example Ljung [7]. We can, therefore, use standard system identification techniques to identify the coefficients of Equation (2.25) and simply rearrange the equation to obtain a parity relation. Seen in another way, constructing a robust parity relation is equivalent to finding a robust ARX model for the plant.

2.2 Double Sensor Parity Relations

In some practical cases single sensor parity relations do not provide a reliable indication of sensor failures. By using combinations of two or more sensors it is possible to construct more complex parity relations. The different combinations must be selected so that it would still be possible to identify which sensor has failed. One such method, which will be referred to as double sensor parity relations (DSPR), combines the outputs of two sensors. The double sensor parity relations are derived as follows: let the i th and j th measurements be given by

$$y_i(kT_s) = \mathbf{c}'_i \mathbf{x}(kT_s) + \mathbf{d}'_i \mathbf{u}(kT_s), \quad (2.26)$$

$$y_j(kT_s) = \mathbf{c}'_j \mathbf{x}(kT_s) + \mathbf{d}'_j \mathbf{u}(kT_s), \quad (2.27)$$

where c'_i , c'_j , d'_i and d'_j are the i th and j th rows of C and D respectively. Similar to the single sensor case, we write down a set of equations that relates consecutive outputs with an initial state and the inputs to the system:

$$\begin{aligned}
y_i(k) &= c'_i x(k) + d'_i u(k), \\
y_j(k) &= c'_j x(k) + d'_j u(k), \\
y_i(k+1) &= c'_i A x(k) + c'_i B u(k) + d'_i u(k+1), \\
y_j(k+1) &= c'_j A x(k) + c'_j B u(k) + d'_j u(k+1), \\
&\vdots \\
y_i(k+n_i-1) &= c'_i A^{n_i-1} x(k) + c'_i A^{n_i-2} B u(k) + \dots + d'_i u(k+n_i-1), \\
y_j(k+n_j) &= c'_j A^{n_j} x(k) + c'_j A^{n_j-1} B u(k) + \dots + d'_j u(k+n_j), \\
y_i(k+n_i) &= c'_i A^{n_i} x(k) + c'_i A^{n_i-1} B u(k) + \dots + d'_i u(k+n_i), \tag{2.28}
\end{aligned}$$

where we assume that $n_i = n_j + 1$. These equations can again be written in a more compact form similar to Equation (2.13) but, to simplify notation, we will first reorder the equations so that all the equations involving y_i appear first. We then have

$$\begin{pmatrix} y_i(n_i) \\ y_j(n_j) \end{pmatrix} = C_{ij} x(k) + D_{ij} u(n_i), \tag{2.29}$$

where

$$y_i(n_i) = [y_i(k); y_i(k+1); \dots; y_i(k+n_i)], \tag{2.30}$$

$$y_j(n_j) = [y_j(k); y_j(k+1); \dots; y_j(k+n_j)], \tag{2.31}$$

$$C_{ij} = [c'_i; c'_i A; \dots; c'_i A^{n_i}; c'_j; c'_j A; \dots; c'_j A^{n_j}], \tag{2.32}$$

$$D_{ij} = \begin{pmatrix} D_i \\ D_j \end{pmatrix}, \quad (2.33)$$

where D_i and D_j are defined by Equation (2.17) with $n_i + 1$ and $n_j + 1$ rows respectively. Because we have assumed that n_j is one less than n_i , the last n_u columns of D_j will be zero because $y_j(k + n_j)$ does not depend on $\mathbf{u}(k + n_i)$. The condition for constructing a double sensor parity relation is given by Chow [4]: the observable subspaces of the i th and j th sensors must overlap. Assuming this is the case, we can find vectors β_i and β_j so that

$$[\beta_i^T, \beta_j^T] C_{ij} = 0. \quad (2.34)$$

Multiplying Equation (2.29) with $[\beta_i^T, \beta_j^T]$ we get the ij th double sensor parity relation

$$\sum_{s=0}^{n_i} \beta_{i,s} y_i(k-s) + \sum_{s=1}^{n_j} \beta_{j,s} y_j(k-s) - \sum_{r=1}^{n_u} \sum_{s=0}^{n_i} \alpha_{ij,r,s} u_r(k-s) = 0, \quad (2.35)$$

where

$$\alpha_{ij} = [\beta_i^T, \beta_j^T] D_{ij}. \quad (2.36)$$

A block diagram of the DSPR Residual Generator is shown Figure 2.2. If either the i th or




Figure 2.2: Block diagram of DSPR Residual Generator.

the j th sensor fails the above relations will not hold; we define the ij th DSPR residual r_{ij}

as

$$r_{ij}(k) = \sum_{s=0}^{n_i} \beta_{i,s} y_i(k-s) + \sum_{s=1}^{n_j} \beta_{j,s} y_j(k-s) - \sum_{r=1}^{n_u} \sum_{s=0}^{n_i} \alpha_{ij,r,s} u_r(k-s). \quad (2.37)$$

In general, when the i th sensor fails, the set of residuals r_{iq} , $i < q \leq n_y$ and r_{pi} , $1 \leq p < i$ will all be nonzero. This set uniquely identifies the i th sensor.

If, instead of using the i th measurement as the last row in Equation (2.28) we use the j th measurement, n_j will equal $n_i + 1$ and we get a dual relation and residual. We will refer to these as the j th DSPR and residual respectively. The residual in this case is

$$r_{ji}(k) = \sum_{s=1}^{n_i} \beta_{i,s} y_i(k-s) + \sum_{s=0}^{n_j} \beta_{j,s} y_j(k-s) - \sum_{r=1}^{n_u} \sum_{s=0}^{n_j} \alpha_{ji,r,s} u_r(k-s). \quad (2.38)$$

2.3 Actuator Parity Relations

In the example at the beginning of this chapter we have shown that generalized parity relations can be used to detect actuator failures. Dutilloy [5] has shown how to construct actuator parity relations given the discrete-time system description, Equations (2.8) and (2.9), for the case $D = 0$. The case where D is nonsingular will be treated here. To construct the actuator parity relations we again find the output history as in Equation (2.12) but now we must use the same number of sensors as actuators, i.e., we must use a subset of sensors so that $n_y = n_u$. The reason for this requirement will become clear later in the derivation. We will assume that this is the case and that the output is given by Equation (2.9). The set of output equations can be written as a matrix-vector equation

$$y(n_i) = Cx(k) + Du(n_i), \quad (2.39)$$

where

$$y(n_i) = [y(k); y(k+1); \dots; y(k+n_i)], \quad (2.40)$$

$$u(n_i) = [u(k); u(k+1); \dots; u(k+n_i)] \quad (2.41)$$

$$= [u_1(k); u_2(k); \dots; u_{n_u}(k); u_1(k+1); u_2(k+1); \dots; u_{n_u}(k+1); \dots; u_1(k+n_i); u_2(k+n_i); \dots; u_{n_u}(k+n_i)], \quad (2.42)$$

$$C = [C; CA; \dots; CA^{n_i}], \quad (2.43)$$

$$D = \begin{pmatrix} D & 0 & 0 & \dots & 0 \\ CB & D & 0 & \dots & 0 \\ CAB & CB & D & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ CA^{n_i-1}B & CA^{n_i-2}B & CA^{n_i-3}B & \dots & D \end{pmatrix}, \quad (2.44)$$

$y \in \mathbb{R}^{(n_i+1)n_y}$, $u \in \mathbb{R}^{(n_i+1)n_u}$, $C \in \mathbb{R}^{(n_i+1)n_y \times n_x}$, and $D \in \mathbb{R}^{(n_i+1)n_y \times (n_i+1)n_u}$. Because we have chosen $n_y = n_u$ the matrix D will be square. Assuming D is invertible, we can multiply Equation (2.39) by D^{-1} and after rearranging we get

$$u(n_i) = (-D^{-1}C)x(k) + D^{-1}y(n_i). \quad (2.45)$$

This equation is similar to Equation (2.13) with the roles of the outputs and the inputs interchanged. By proceeding as before, we can construct single actuator parity relations (SAPR) and double actuator parity relations (DAPR). A little more work is necessary for the actuator case because $u(n_i)$ contains all the elements of the input in an interleaved way as shown in Equation (2.42). For example, if we want to construct a SAPR for the i th actuator, we must form a vector of inputs that has only u_i 's as elements, starting with $u_i(k)$ and taking every n_u th element of $u(n_i)$. In order to refer to the rows of $D^{-1}C$ and D^{-1} in an easy way we define the following temporary matrices

$$\widetilde{C} = -D^{-1}C \quad (2.46)$$

$$= [\tilde{c}'_1; \dots; \tilde{c}'_{(n_i+1)n_u}], \quad (2.47)$$

$$\widetilde{D} = D^{-1} \quad (2.48)$$

$$= [\tilde{d}'_1; \dots; \tilde{d}'_{(n_i+1)n_u}]. \quad (2.49)$$

We can now set up equations similar to Equation (2.12) for the i th actuator,

$$\bar{u}_i(n_i) = [u_i(k); u_i(k+1); \dots, u_i(k+n_i)] \quad (2.50)$$

$$= \bar{C}_i x(k) + \bar{D}_i y(n_i), \quad (2.51)$$

$$\bar{C}_i = [\tilde{c}'_i; \tilde{c}'_{i+n_u}; \dots; \tilde{c}'_{i+n_u n_i}] \in \mathbb{R}^{(n_i+1) \times n_x}, \quad (2.52)$$

$$\bar{D}_i = [\tilde{d}'_i; \tilde{d}'_{i+n_u}; \dots; \tilde{d}'_{i+n_u n_i}] \in \mathbb{R}^{(n_i+1) \times (n_i+1)n_u}. \quad (2.53)$$

We now find a vector α_i so that

$$\alpha_i^T \bar{C}_i = 0. \quad (2.54)$$

The i th SAPR residual is defined as

$$r_i(k) = \alpha_i^T \bar{D}_i y(n_i) - \alpha_i^T \bar{u}_i(n_i) \quad (2.55)$$

$$= \beta_i^T y(n_i) - \alpha_i^T \bar{u}_i(n_i) \quad (2.56)$$

$$= \sum_{r=1}^{n_y} \sum_{s=0}^{n_i} \beta_{i,r,s} y_r(k-s) - \sum_{s=0}^{n_i} \alpha_{i,s} u_i(k-s), \quad (2.57)$$

where

$$\beta_i^T = \alpha_i^T \bar{D}_i \quad (2.58)$$

$$= [\beta_{i,1,n_i}; \beta_{i,2,n_i}; \dots; \beta_{i,n_u,n_i}; \beta_{i,1,n_i-1}; \beta_{i,2,n_i-1}; \dots; \beta_{i,n_u,n_i-1}; \dots; \beta_{i,1,0}; \beta_{i,2,0}; \dots; \beta_{i,n_u,0}], \quad (2.59)$$

Because of the requirement that $n_y = n_u$, it was found that there is usually more than one vector in the left null space of \bar{C}_i . These vectors give true parity relations (see Lou

et al. [8]) as they all satisfy Equation (2.54) exactly. It is not clear at this point how to select between the different vectors, and whether one is necessarily “better” than another. A block diagram of the SAPR Residual Generator is shown in Figure 2.3.




Figure 2.3: Block diagram of SAPR Residual Generator.

In a similar way we can construct DAPR of the form r_{ij} and r_{ji} . Although we will show experimental DAPR results, we will not derive the equations here as the procedure leading to the results is analogous to the single actuator case.

2.4 Example

To illustrate some of the ideas discussed in the foregoing sections, we present a simple example of a second order system. Many practical systems, including the Mini-Mast which we will discuss in more detail later, are described by the following m -mode state-space model

$$\dot{\mathbf{x}}(t) = \begin{pmatrix} \mathbf{A}_1 & \mathbf{O} \\ & \ddots \\ \mathbf{O} & \mathbf{A}_m \end{pmatrix} \mathbf{x}(t) + \begin{pmatrix} \mathbf{B}_1 \\ \vdots \\ \mathbf{B}_m \end{pmatrix} \mathbf{u}(t) \quad (2.60)$$

where

$$\mathbf{A}_i = \begin{pmatrix} 0 & 1 \\ -\omega_i^2 & -2\zeta_i\omega_i \end{pmatrix}, \quad i = 1, \dots, m, \quad (2.61)$$

$$\mathbf{B}_i = \begin{pmatrix} 0 & \dots & 0 \\ b_{i,1} & \dots & b_{i,n_u} \end{pmatrix}, \quad i = 1, \dots, m, \quad (2.62)$$

where ω_i is the natural frequency of the i th mode with corresponding damping ratio ζ_i . We will analyze only one of the second order blocks. In order to simplify some of the hand calculations we will further write the continuous-time state-space model in the observable canonical form (see Chen [3])

$$\dot{\mathbf{x}}(t) = \begin{pmatrix} 0 & -\omega_n^2 \\ 1 & -2\zeta\omega_n \end{pmatrix} \mathbf{x}(t) + \begin{pmatrix} \omega_n^2 \\ 0 \end{pmatrix} u(t), \quad (2.63)$$

$$y(t) = [0 \quad 1] \mathbf{x}(t) \quad (2.64)$$

$$= \mathbf{c}' \mathbf{x}(t). \quad (2.65)$$

The following parameters will be used:

sampling period $T_s = 0.015$ seconds,

natural frequency $\omega_n = 5$ rad/s (0.8 Hz),

damping ratio $\zeta = 0.01$.

The discretized system is given by

$$\mathbf{x}(k+1) = \begin{pmatrix} 0.9972 & -0.3774 \\ 0.0150 & 0.9957 \end{pmatrix} \mathbf{x}(k) + \begin{pmatrix} 0.3746 \\ 0.0028 \end{pmatrix} u(k)$$

$$= \mathbf{A}\mathbf{x}(k) + \mathbf{b}u(k), \quad (2.66)$$

$$\begin{aligned} y(k) &= [0 \quad 1]\mathbf{x}(k) \\ &= \mathbf{c}'\mathbf{x}(k). \end{aligned} \quad (2.67)$$

We can also write this single-input single-output system as a difference equation

$$\begin{aligned} y(z) &= \mathbf{c}'(z\mathbf{I} - \mathbf{A})^{-1}\mathbf{b}u(z) \\ &= \frac{n(z)}{d(z)}u(z) \\ &= \frac{b_{21}z^{-1} + (a_{21}b_{11} - a_{11}b_{21})z^{-2}}{1 - (a_{11} + a_{22})z^{-1} + (a_{11}a_{22} - a_{12}a_{21})z^{-2}}u(z) \\ &= \frac{0.002810z^{-1} + 0.002808z^{-2}}{1 - 1.992883z^{-1} + 0.998501z^{-2}}u(z). \end{aligned} \quad (2.68) \quad (2.69)$$

The difference equation describing the system is

$$y(k) - 1.992883y(k-1) + 0.998501y(k-2) = 0.002810u(k-1) + 0.002808u(k-2). \quad (2.70)$$

The SSPR residual is easily found as

$$r = \sum_{s=0}^2 \beta_s y(k-s) - \sum_{s=1}^2 \alpha_s u(k-s), \quad (2.71)$$

where

$$\boldsymbol{\beta} = [0.998501; \quad -1.992883; \quad 1], \quad (2.72)$$

$$\boldsymbol{\alpha} = [0.002808; \quad 0.002810; \quad 0]. \quad (2.73)$$

Note that $\alpha_0 = 0$; this is expected because there is no direct feedforward from the input to the output. The plant and Residual Generator are shown schematically in Figure 2.4. Note that the transfer functions of the Residual Generator are the numerator and denominator of the transfer function of the plant — the residual is formed by multiplying the output $y(z)$ by the denominator polynomial, the input $u(z)$ by the numerator polynomial, and

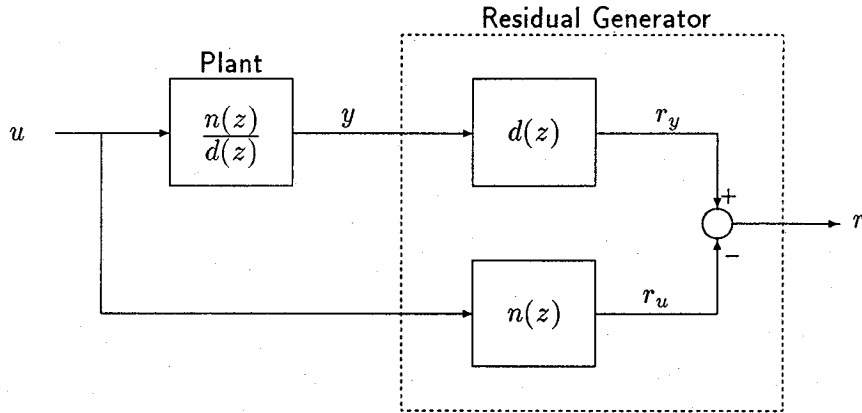


Figure 2.4: Block diagram of the plant and SSPR Residual Generator.

subtracting the latter from the former. The transfer functions for this Residual Generator are shown in Figure 2.5. The transfer function from y to r has a large magnitude at high frequencies. This will always be the case for practical systems as they have a natural roll-off at high frequencies. The high gain at high frequencies can be a source of trouble if we have noisy sensors or unmodelled high frequency dynamics.

The coefficients multiplying the input sequence are very small — it was first believed that this is due to the small damping in the system but it is easily shown that this is not necessarily the case. By repeating the above example and changing the damping ratio by a factor of ten to $\zeta = 0.1$, we get the following coefficients:

ζ	α_2	α_1
0.01	0.002808	0.002810
0.10	0.002783	0.002797

The discretization step was also carried out symbolically and the detail can be found in




Figure 2.5: Transfer functions of the SSPR Residual Generator. The transfer functions are periodic and are shown up to half the sampling frequency.

Appendix A. We see that the elements of the \mathbf{A} and \mathbf{B} matrices have factors like $e^{-\zeta\omega_n T_s}$, $\cos(\omega_n\sqrt{1-\zeta^2}T_s)$ and $\sin(\omega_n\sqrt{1-\zeta^2}T_s)$. The small coefficients are a result of the product of ζ , ω_n , and T_s . Even if we had a larger damping ratio ζ , these elements of α will still be small because T_s is small. For a given practical system we have no control over ζ and the only parameter that we can vary (to a limited degree) is the sampling period.

For the single-input single-output case, the single actuator parity relation is identical to the single sensor parity relation. Therefore, only one relation exists and it is not possible to

determine from a nonzero residual alone whether it was an actuator or sensor failure.

Chapter 3

Displacement Sensor Failure Detection

3.1 Introduction

In this chapter we discuss a series of failure detection experiments that were conducted on the Mini-Mast. Specifically, we will look at the detection of displacement sensor failures of the Mini-Mast and discuss several factors that influence the performance of the Residual Generators. We will also compare parity relations obtained from a state-space model with parity relations identified directly on a set of input-output data. The parity relations obtained from the state-space model will be referred to as the model-based relations and those obtained by identification as the identified relations. First, we give a brief description of the Mini-Mast.

The Mini-Mast is an experimental truss at the NASA Langley Research Center, Hampton,

Virginia. The mast is deployed vertically and is rigidly fixed at its base. It has 18 bays, each of length 1.12 meter (3.68ft); the total length of the mast is 20.16 meters (66.14ft). The bays are numbered 1 through 18, with Bay 18 at the top. The mast has three member types: longerons, battens, and diagonals. Longerons are parallel to the vertical axis and provide beam stiffness and strength in bending. Battens are in the beam face planes and provide stability. Diagonals, also in the beam face planes, provide stiffness and strength in torsion and shear. The mast is shown schematically in Figure 3.1. The truss has 57 corner joints with stainless steel pins that allow the longerons and diagonal members to be hinged, so that it is possible to retract and deploy the mast. Three torque wheel actuators




Figure 3.1: Schematic diagram of the Mini-Mast and orientation of the displacement sensors. The sensors measure displacements normal to their surfaces, relative to a fixed structure.

are mounted at the top of the mast parallel to the XYZ axes. By applying voltages to these motors, it is possible to apply torsional and bending torques to the mast. These actuators were used in the failure detection experiments to excite the mast. The mast is also instrumented with a full set of accelerometers, rate gyros, and displacement sensors. The displacement sensors are mounted so that each measures displacements normal to its reference surface, and relative to a fixed structure that is built around the mast. Three displacement sensors are mounted at each bay but only the three sensors at Bay 18 were used.

A finite element model for the Mini-Mast has been developed by NASA to analyze the modal frequencies and mode shapes. A brief summary is given here; detail can be found in Pappa et al. [13]. The first two modes are the first bending modes, oriented in the X and Y directions. The natural frequencies of these modes are approximately 0.65Hz. This is followed by the first torsion mode with a natural frequency of approximately 4.4Hz. The fourth and fifth modes are the second bending modes with natural frequencies of approximately 6.2Hz. The directions of the second bending modes are rotated by 45 degrees from the X - Y directions, thus coupling the bending responses. The first and second of 108 local modes, caused mainly by the diagonal members, have natural frequencies of approximately 14.8Hz. Other modes are: second torsional at 20.86Hz, third bending modes at 29.79Hz and 30.94Hz, third torsional at 38.83Hz, fourth bending modes at 40.12Hz and 43.41Hz, fourth torsional at 54.30Hz, fifth bending modes at 66.34Hz and 70.25Hz, and fifth torsional mode at 71.88Hz. The state-space model used to generate the model-based parity relations included the first 5 modes of the system; the modal frequencies and damping ratios used are shown in Table I. The state-space model was obtained by Drs. Raymond Montgomery and David Ghosh of NASA Langley Research Center by an analysis of input-output data in preparation for the design of a control system for the Mini-Mast. The state-space matrices are given in

Appendix B.

Table I. State-space model modal frequencies and damping ratios

Mode	ζ	ω [Hz]	ω [rad/s]
First bending	0.0323	0.8559	5.3778
First bending	0.0213	0.8547	5.3702
First torsional	0.0717	4.2933	27.0133
Second bending	0.0238	6.1186	38.4440
Second bending	0.0100	6.1669	38.7478

Several experiments were conducted on the Mini-Mast to obtain input-output data sets. The mast was excited by driving the torque wheels with random signals. For the experiments discussed in this chapter, the input signal amplitudes were independent, identically distributed with a uniform probability density function. The sampling period T_s was 15 ms. This is a baseline sampling period that will be used by the control system for the mast. The input signals were held constant for four sampling periods, i.e., for 60 ms. This choice gave the freedom to simulate different sampling periods when analyzing the sensor parity relations. Unfortunately, keeping the amplitude constant for more than one sampling period but taking samples every sampling period results in a signal with a spectrum that has zeros at frequencies lower than half the sampling frequency. A typical spectrum of an input signal that was held constant for four sampling periods but that was sampled every sampling period is shown in Figure 3.2. Fortunately, due to nonlinearities of the actuators and joints of the Mini-Mast, no zeros occurred in the output spectrum.




Figure 3.2: Spectrum of the input signal. The input was held constant for 4 sampling periods ($4T_s$) but samples were taken every sampling period, $T_s = 15$ ms.

The three displacement sensors at the tip of the mast will be referred to as Sensor D1, Sensor D2 and Sensor D3 with corresponding measurements y_1 , y_2 and y_3 and SSPR residuals r_1 , r_2 and r_3 . The transfer functions from the i th measurement y_i to the i th residual r_i will be called $B_i(z)$ and the transfer functions from the inputs u_1, \dots, u_{n_u} to r_i will be denoted by $A_{i,1}(z), \dots, A_{i,n_u}(z)$. In some experiments we will use an increased sampling period of 30 ms, which is twice the baseline sampling period; this will be referred to as $2T_s$. The order of the parity relation, n_i in Equation (2.25), will be referred to as the number of lags. Note that for n_i lags we are actually using $n_i + 1$ samples of the corresponding measurement: n_i past values plus the current sample. Corresponding to the 10 dimensional state of the state-space model used, the model-based parity relations incorporate 10 lags.

The spectrum of y_2 is shown in Figure 3.3. In this figure we clearly see the first bending mode




Figure 3.3: Spectrum of Displacement Sensor 2.

at approximately 0.9Hz and the first torsional mode at 4.3Hz. The peaks in the spectrum at 12.6Hz, 13.9Hz and 16.6Hz correspond to the local modes. The second torsional mode is at approximately 21.4Hz. Further, though the input signals have zeros in their spectra (see Figure 3.2), they do not show up in the spectrum of the output signal. Note that 256 point DFTs were used to compute these spectra so that we do not have very fine spectral resolution. The spectra of the other two displacement sensors are similar in nature to the one just shown and will not be shown here. When we refer to a particular behavior of a residual later in this work only one example will be given to illustrate the point. If a specific example does not represent all the sensors it will be noted explicitly.

Failures of the sensors were simulated in the data by modifying the recorded data. In most of the examples that we will discuss the sensor is failed to zero by simply zeroing the output data. (See Equation (2.7) for the modelling of failures.) We will also choose the failure times to be approximately in the middle of a plot so that it will be easy to compare the amplitude of the residual before and after the failure.

3.2 Model-based Single Sensor Parity Relations

Figure 3.4 shows the failure of Sensor D1 that has failed to zero at sample number 213. The failure is clearly indicated by the large transient in the residual. In this figure we also see a behavior that was typical for all model-based residuals for displacement sensors: the residual has a large amplitude while the sensor is in perfect condition followed by a smaller amplitude when the transients excited by the failure are gone. In Chapter 2 it was shown that the inputs to the i th Residual Generator are all the control inputs and, for single sensor parity relations, the i th measurement. Equation (2.22) further shows that the i th residual r_i has two components $r_{i,y}$ and $r_{i,u}$, corresponding to the i th measurement and all the inputs. The residual is defined as the difference between these two components. Therefore, except for noise and unmodelled effects, we expect these two components to be equal. Plotting the components $r_{1,y}$ and $r_{1,u}$ separately in Figure 3.5, we see that this is not so. The component $r_{1,y}$ has a much larger amplitude than $r_{1,u}$ and there is no similarity between the two components. At first it was believed that this discrepancy is due to the small damping of the mast but the example at the end of Chapter 2 clearly indicates that this is not the reason. This difference in amplitude of the two components explains the previously mentioned behavior that the residual amplitude is large while the sensor is fully operational and small when the sensor has failed. The reason for the mismatch will be given




Figure 3.4: Displacement Sensor D1 failure. Top: Sensor D1 output y_1 . Bottom: model-based SSPR residual r_1 . Sensor D1 has failed to zero at sample number 213.

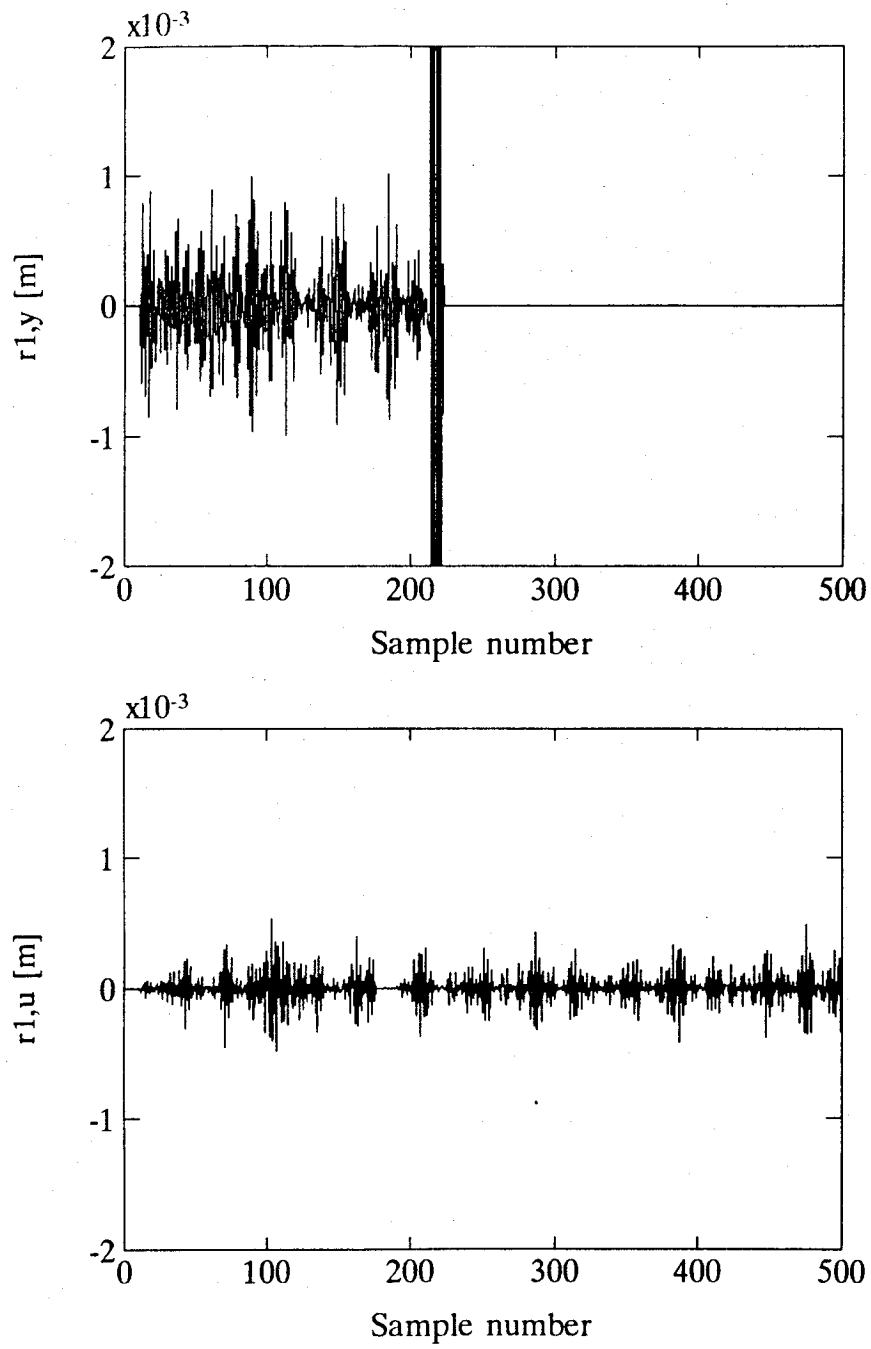


Figure 3.5: Components $r_{1,y}$ (top) and $r_{1,u}$ (bottom) of model-based SSPR r_1 . Sensor D1 has failed to zero at sample number 213.

when we discuss the transfer functions of the Residual Generator.

The SSPR residual r_3 is shown in the top of Figure 3.6. In this example Sensor D3 has failed to zero at sample number 235. As before, we see a large transient when the failure occurs. The bottom of Figure 3.6 shows the same residual, but this time Sensor D3 has failed at sample number 234, one sample (15 ms) earlier. Although a brief pulse is visible, we did not get a clear failure signature and the spike could have been caused by noise. This inability of the model-based single sensor parity relations to give a clear indication of sensor-off failure modes occurred often and the reason for the poor performance will be explained later. We now show a different failure mode.

A noisy sensor was simulated by adding white noise to the output of Sensor D2. The plot at the top of Figure 3.7 shows the output of Sensor D2 with noise added to it from sample number 240. The standard deviation of the noise was one hundredth that of the standard deviation of the measurement y_2 . The effect of the noise is barely visible in the measurement. The corresponding SSPR residual, r_2 , is shown in the bottom of Figure 3.7. The failure is clearly indicated by the residual. So the added-noise failure mode is clearly detected by the parity relation. However, this extreme sensitivity of the residual to noise can be a problem when we are working in a really noisy environment. Before we discuss the transfer functions of the Residual Generators we first turn to parity relations identified on input-output data.

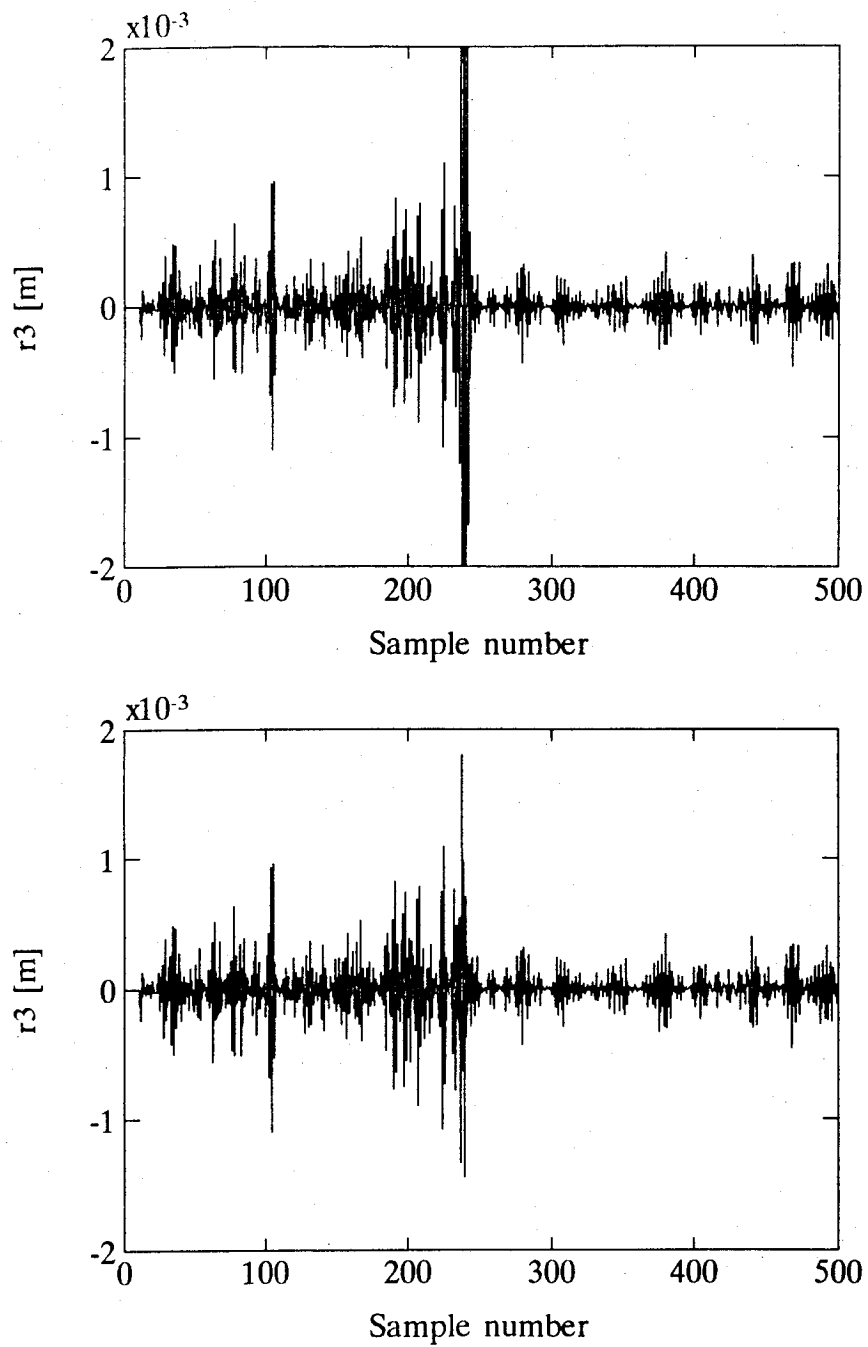


Figure 3.6: Top: model-based SSPR r_3 when Sensor D3 has failed to zero at sample number 235. Bottom: the same residual when Sensor D3 failed at sample number 234, one sample earlier.





Figure 3.7: Top: Sensor D2 output. Noise was added to Sensor D2 from sample number 240. Bottom: Model-based SSPR r_2 .

3.3 Identified Single Sensor Parity Relations

It was noted in Chapter 2 that single sensor parity relations correspond to an ARX model of the plant. Using a different set of input-output data, the coefficients of the parity relation (see Equation (2.25)) were identified using a least squares criterion. The length of the data set was slightly less than 30 seconds. These parity relations, which will be referred to as identified relations, were applied to the same data used in Section 3.2. Figure 3.8 shows




Figure 3.8: Identified SSPR residual r_3 . Sensor D3 has failed to zero at sample number 234. Compare with the plot at the bottom of Figure 3.6.

the identified SSPR residual r_3 when Sensor D3 has failed to zero at sample number 234, i.e., at the same time as portrayed in the bottom graph of Figure 3.6. In that case the model-based SSPR failed to give a clear indication of the failure. In Figure 3.8 we see that the identified residual gives a very different failure signature. First, note that the

amplitude of the identified residual is smaller than the amplitude of the model-based residual by approximately two orders of magnitude. Furthermore, the amplitude of the identified residual is small while the sensor is in good condition and large while the sensor is faulty, the opposite of what we had before. Clearly, this case is much closer to what we would like to see. To highlight the difference between the model-based and identified relations, we show the components $r_{3,y}$ and $r_{3,u}$ in Figure 3.9. Here we see that the contributions $r_{3,y}$ and $r_{3,u}$ are approximately of the same magnitude. We also see in these figures that the two components have similar wave forms and thus, when subtracted from each other, will result in a residual with a small amplitude. Careful comparison between Figures 3.6 and 3.9 further shows that, while the sensor is in working condition, the model-based residual has more high frequency content than the identified residual. The reason for this will become clear when we discuss the different Residual Generator transfer functions in the next section.

With the identified relations we have the luxury of easily increasing the number of lags used in the parity relations. In Figure 3.10 we show the residual of an identified SSPR relation with 20 lags. To make a comparison with a previous failure we have chosen a failure of Sensor D3 at sample number 234. Comparing Figure 3.10 with Figure 3.8 we see that increasing the number of lags results in a residual with a smaller amplitude while the sensor is in good health and a slightly larger residual when the failure is present. Therefore, at the expense of an increase in the number of computations, we can improve the failure signature by choosing a higher order model.





Figure 3.9: Components $r_{3,y}$ and $r_{3,u}$ of identified SSPR residual r_3 .




Figure 3.10: Identified SSPR residual r_3 with 20 lags. Sensor D3 has failed to zero at sample number 234.

3.4 Transfer functions of model-based and identified Single Sensor Parity Relations

To explain some effects that we have seen in the preceding sections and further highlight the differences between the model-based and identified SSPR residuals, we now turn to the transfer functions of the corresponding residual generators.

In Chapter 2 it was noted that a SSPR Residual Generator is a multi-input single-output finite impulse response filter so that the individual transfer functions have no poles (except for poles at the origin). The zero locations of the model-based and identified Residual Generators for the transfer function $B_2(z)$ are shown in Figure 3.11. We see that the identified relation has zeros at higher frequencies than the model-based relation. The zeros of the model-based Residual Generator are simply the poles of the plant (see Section 2.1), and the poles have been constrained to the first five modes of the mast by our selection of the model. During the identification process no constraint is placed on the pole locations and the resulting model thus gives poles that give the best fit over all frequencies. Except for one complex zero pair, there is little correspondence between the zero locations of the two transfer functions.

The transfer functions of the model-based and identified Residual Generators are compared in Figure 3.12. We first note that the model-based transfer function from y_2 to r_2 has a small gain at low frequencies and a high gain at high frequencies. This high gain at high frequencies explains the extreme sensitivity that the residual showed to a noisy sensor (see Figure 3.7). Although it was not shown there, the corresponding identified residual was less sensitive to noise. The high gain is also responsible for the good transient that we have seen in Figure 3.4. In that figure we see that there was an abrupt change in the measurement at the time of failure. For the example shown at the bottom of Figure 3.6, the time of




Figure 3.11: Left: Zero locations of the model-based Residual Generator transfer function y_2 to r_2 . Right: Zero locations of the identified Residual Generator transfer function y_2 to r_2 . The solid line circles have radius 1. Note that the model-based Residual Generator has two closely spaced zeros at approximately 45 degrees.

failure was chosen so that the output y_3 was close to a zero crossing point so that there was no abrupt change in the signal. The high gain at high frequencies also explains why the components $r_{2,y}$ and $r_{2,u}$ have such different amplitudes — noise in the measurement is amplified considerably so that the contribution of that component is much larger than the contribution of u_1, \dots, u_{n_u} . The model-based transfer functions $A_{2,1}(z), \dots, A_{2,3}(z)$ also have smaller gains at low frequencies than the identified relations. The identified relation clearly puts more emphasis at low frequencies and less at high frequencies.

The spectra of the model-based and identified residuals are shown in Figure 3.13. We see that the model-based residual has very little frequency content at low frequencies and much greater frequency content at high frequencies. Note that the difference of the minimum

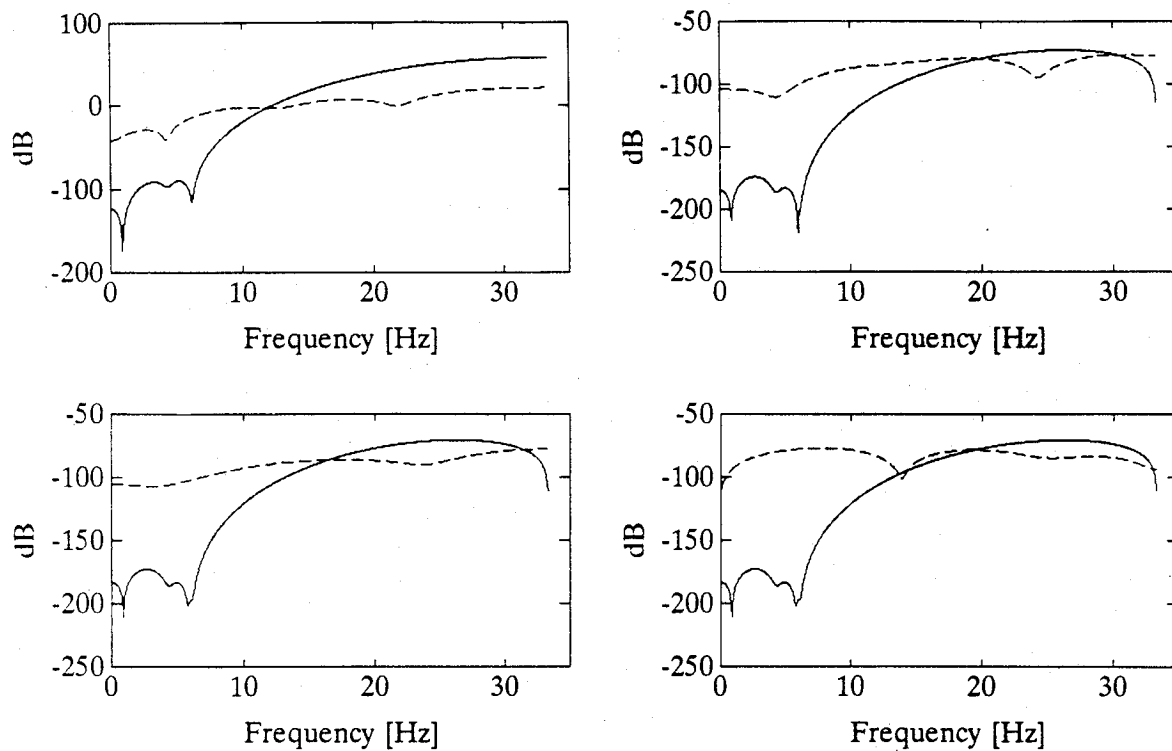


Figure 3.12: Transfer functions of the model-based (solid line) and identified (dashed line) SSPR Residual Generators. Top left: y_2 to r_2 , top right: u_1 to r_2 , bottom left: u_2 to r_2 , bottom right: u_3 to r_2 .

at low frequencies and the maximum at high frequencies is almost 180 dB! Clearly, the model-based Residual Generator does a very good job at frequencies below 7 Hz. However, because we have a reduced order model with an excellent match at low frequencies, there is a significant mismatch at high frequencies and this prevents the model-based relations from obtaining good performance. The large high frequency content was pointed to earlier when we discussed the differences between model-based and identified relations in Figures 3.5, 3.6 and 3.9. Note further that the model-based spectrum clearly shows a peak at approximately 14.4 Hz that corresponds to the local modes which are not included in the state-space model.




Figure 3.13: Spectra of the model-based and identified residuals.

The spectra of the identified residuals exhibit an almost flat response over all frequencies. The limitation of the 10 lag relation clearly shows up as a peak at approximately 0.9 Hz, the first bending mode, as well as a peak at approximately 6 Hz, the second bending mode. Increasing the number of lags to 20 clearly shows an improved match at the first bending mode and a spectrum with a slightly smaller magnitude over most of the frequency band.

In the next section we investigate the effect of the sampling period on the performance of the Residual Generators.

3.5 Increased Sampling Period

It was found that increasing the sampling period had a significant effect on the identified SSPR residuals. An increase in the sampling period gave improvement on the model-based SSPRs. Using the same data set as before the sampling period was increased to 30 ms, i.e., $2T_s$. Figure 3.14 shows the identified residual r_3 when Sensor D3 has failed to zero at




Figure 3.14: Identified SSPR r_3 , 10 lags, sampling period $2T_s$. Sensor D3 has failed to zero at sample number 117.

sample number 117; this corresponds to the same time as we had in Figures 3.9 and 3.10. Here we clearly see that doubling the sampling period leads to a major improvement in the failure signature. The same failure is shown in Figure 3.15 where we have used a sampling period of 30 ms and a parity relation with 20 lags — an excellent failure signature.




Figure 3.15: Identified SSPR r_3 , 20 lags, sampling period $2T_s$.

It was hoped that the transfer functions of the corresponding Residual Generators would hint at why the increased sampling period leads to so much improvement in the residual but an analysis turned out to be fruitless. One possible reason is that at 10 lags only a small portion of one period of lowest frequency of interest, i.e., the first bending mode at 0.9 Hz, counts in the computation of a relation — with noise contaminating the measurement, it is difficult to capture the underlying low frequency component. Increasing the sampling period results in samples taken further apart so that, using the same number of lags, a greater portion of one period is covered. Another possible reason is that, at $2T_s$, a smaller frequency band needs to be matched by the ARX model leaving more freedom to give a better model at low frequencies. The transfer functions of 20 lag, $1T_s$ and 20 lag, $2T_s$ identified SSPR Residual Generators are compared in Figure 3.16. We see that the $2T_s$ transfer functions tend to have more peaks and dips at low frequencies compared to the $1T_s$

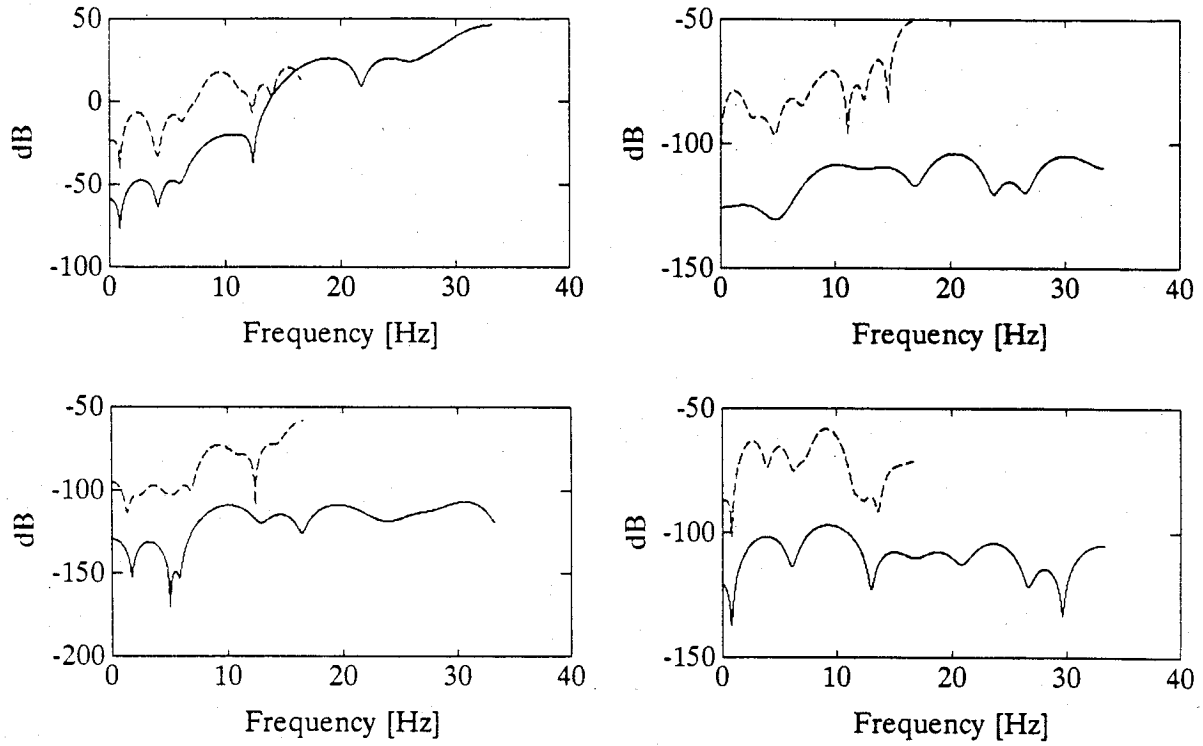


Figure 3.16: Identified SSPR Residual Generator transfer functions, 20 lags, $1T_s$ (solid line) and $2T_s$ (dashed line). The transfer functions are as follows: upper left: y_2 to r_2 , upper right: u_1 to r_2 , lower left: u_2 to r_2 , lower right: u_3 to r_2

counterparts, indicating that more modes are being included at the lower frequencies.

3.6 Double Sensor Parity Relations

In this section we present several failures where DSPRs are used to detect the failure. As before, we will compare model-based relations with identified relations and discuss the effect of increased number of lags and increased sampling period on the performance of the Residual Generators.

Figure 3.17 shows the model-based residuals r_{12} and r_{13} where Sensor D1 has failed to zero at sample number 238. A brief transient is visible at the time of failure. Note further that the residual remains small after the transient is gone. Like the model-based SSPRs, the model-based DSPRs sometimes fail to indicate the failure of a sensor. An example is shown in Figure 3.18 where Sensor D1 has failed to zero at sample number 250. In this example the residuals give no indication of the failures at all. Careful inspection of the plot at the bottom of the figure shows that the first part of the residual up to sample number 250 has a high frequency content while the part from sample number 250 to the end shows some underlying low frequencies. This is to be expected as the DSPR Residual Generator has as inputs the plant inputs u_1 , u_2 , u_3 as well as the two measurements y_1 and y_3 . Therefore, even when Sensor D1 fails to zero, the dynamics of the mast are still being fed to the Residual Generator through Sensor D3. We thus would expect that this signal, which has low frequencies in it, should appear at the output of the Residual Generator.

The detection of the failure of Sensor D2 at sample number 150 by an identified DSPR is shown in Figure 3.19. Both the residuals r_{12} and r_{32} give a clear indication of the failure. The number of lags used was 10. Although this is a different sensor and the parity relations have more lags than the model-based relation, a comparison will still be made. We note that the identified residuals are significantly smaller than the model-based residuals. Furthermore, the difference in frequency content of the residual before and after the failure is large. This invites signal processing to improve the failure signature. It was noted in Section 2.2 that it is possible to construct a dual parity relation for a specific pair of sensors. The dual residuals r_{21} and r_{23} are shown in Figure 3.20. Clearly, there is a marked difference in the amplitudes of the residuals when compared to the ones in Figure 3.19. When this difference was first noted it was believed that this is because Sensor D2 appears as the first sensor in the relation but this big difference did not manifest itself in the other relations

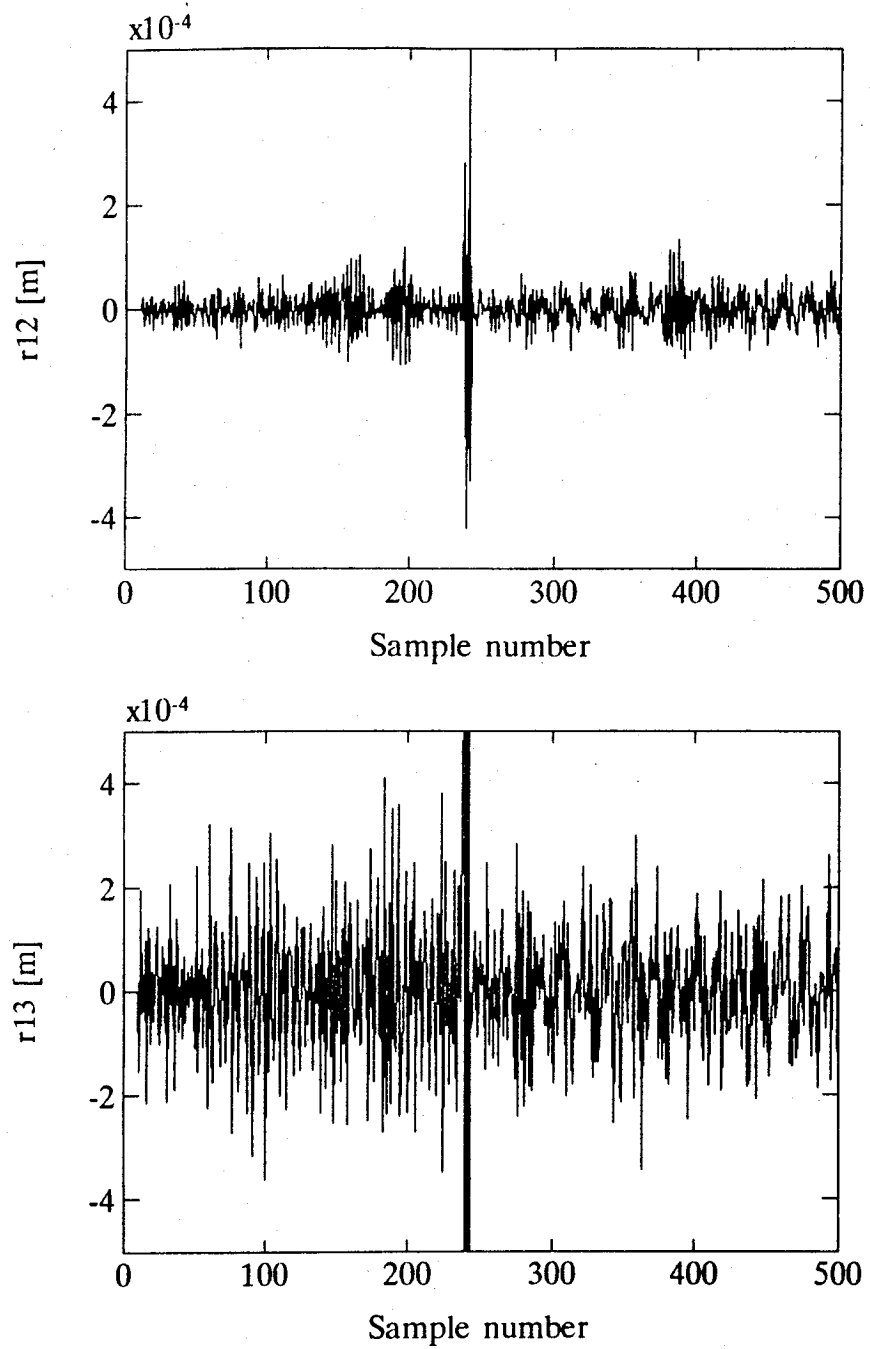


Figure 3.17: Model-based DSPR residuals r_{12} and r_{13} . Sensor D1 has failed to zero at sample number 238.

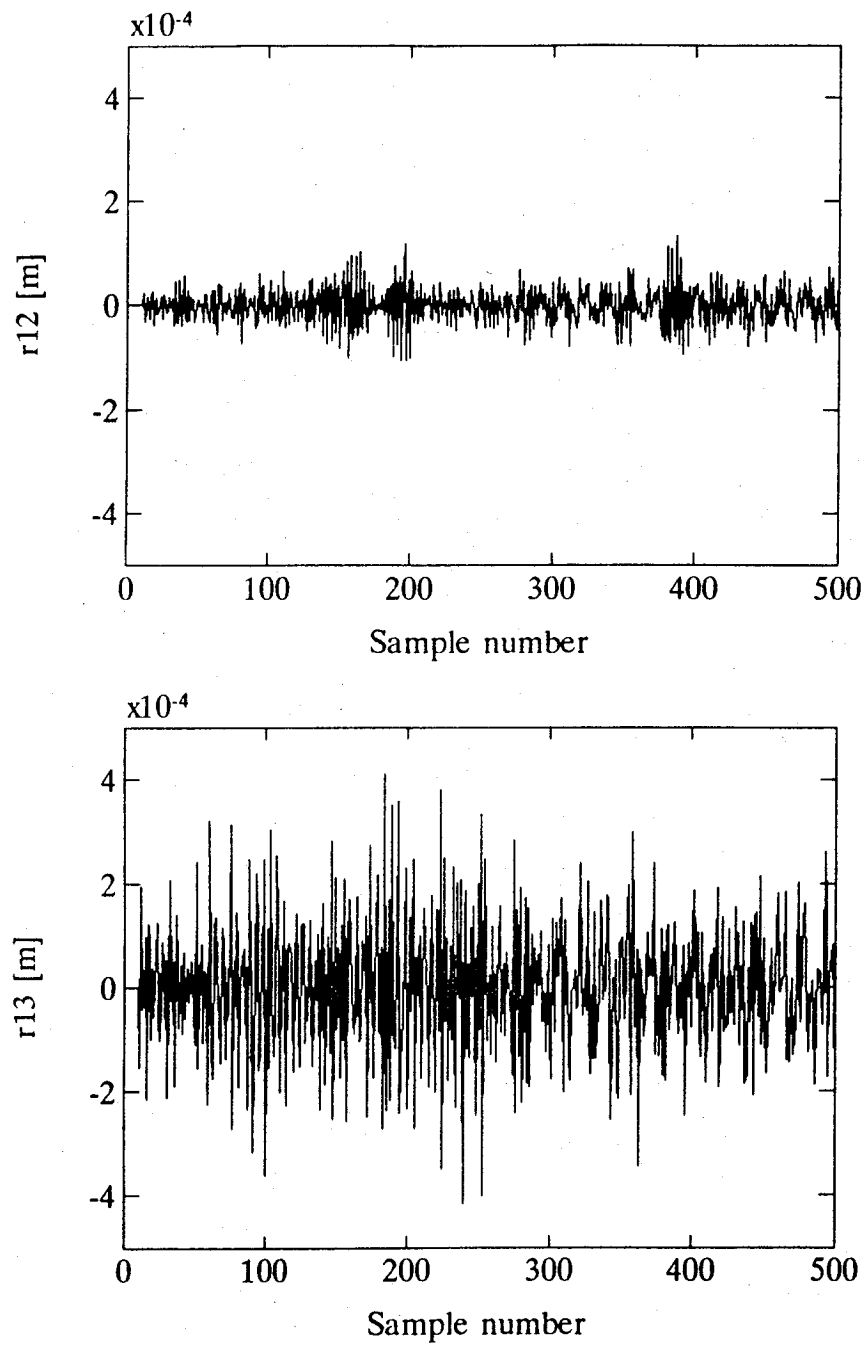


Figure 3.18: Model-based DSPR residuals r_{12} and r_{13} . Sensor D1 has failed to zero at sample number 250.





Figure 3.19: Identified DSPR residuals r_{12} and r_{32} . Sensor D2 has failed to zero at sample number 150.





Figure 3.20: Dual identified DSPR residuals r_{21} and r_{23} . Sensor D2 has failed to zero at sample number 150.

and their dual forms.

Generalized parity relations do not require the specification of the failure mode. The detection of a different type of failure by identified relations is shown in Figure 3.21. Here a gain reduction of 50% in the output of Sensor D3 was simulated from sample number 180 to 500. Both residuals clearly indicate this failure.

Increasing the number of lags in the relations again resulted in improved failure signatures. Figure 3.22 shows the residuals where we have used 20 lags in the DSPRs. This is the same failure that we have seen in Figure 3.20. A comparison of the two figures shows that there is an advantage in increasing the number of lags. The amplitudes of the residuals are smaller when the sensors are in healthy condition and larger once a sensor fails.

Increasing the sampling period again resulted in a significant improvement of the failure signatures. A *model-based* DSPR at $2T_s$ is shown in Figure 3.23. Comparing this figure with Figure 3.18 we notice a significant difference between the residuals. Considering that we are using the same continuous-time state-space model, but now using a longer sampling period, it is clear that the sampling period has a significant effect on the performance of a parity relation. An example of a 20 lag identified DSPR with a $2T_s$ sampling period is shown in Figure 3.24.





Figure 3.21: Identified DSPR residuals r_{13} and r_{23} . The output of Sensor D3 was reduced by a factor of 2 from sample number 180 to 500.





Figure 3.22: Identified DSPR residuals r_{21} and r_{23} , 20 lags. Sensor D2 has failed to zero at sample number 150.





Figure 3.23: Model-based DSPR residuals r_{12} and r_{13} , 10 lags, $2T_s$. Sensor D1 has failed to zero at sample number 125.





Figure 3.24: Identified DSPR residuals r_{21} and r_{23} , 20 lags, $2T_s$. Sensor D2 has failed to zero at sample number 75.

3.7 Summary

In this chapter we have looked at the detection of displacement sensor failures using single and double sensor parity relations. A comparison was made between a set of relations obtained from a state-space model of the Mini-Mast and a set of relations that was obtained by identifying the coefficients of the parity relations directly from a set of input-output data. The state-space model included the first five modes of the mast. The model-based relations failed to indicate all the failures and were very sensitive to noise. The sensitivity to noise is a result of the very large gains at high frequencies of the corresponding Residual Generators. The spectra of the model-based residuals indicate that the state-space model gives an excellent fit at frequencies below 7 Hz at the expense of a poor fit at high frequencies. Reduced order low frequency models are often used in control system design but the results of this chapter show that they are not suitable to design Residual Generators for use in failure detection.

The identified residuals always gave a clear indication of the failure. An analysis of the Residual Generator transfer functions shows that the identified relations put more emphasis at low frequencies and less at the high frequencies. The flat spectra of the residuals suggests that it is important that the model fit the plant well even at high frequencies.

By identifying the parity coefficients directly from input-output data we had the freedom of choosing the model order. In all the experiments an increase in the number of lags (i.e. increasing the order of the model) led to an improvement of the failure signature. An increase in the model order usually resulted in a smaller residual while the sensors were in good health as well as an increase in the magnitude of the residual when a failure was present.

To improve the performance of the Residual Generator in the case of sensor-off failures, double sensors parity relations can be used. In all the experiments and different failure modes considered, the double sensor parity relations performed better than their single sensor parity relation counterparts. The main reason for the improvement is the inclusion of an extra measurement that feeds dynamics of the plant to the Residual Generator even when the other sensor fails to zero.

Increasing the sampling period resulted in a significant improvement of the failure signatures. This is probably because, with a short sampling period, only a small portion of one period of a low frequency is covered by a relation with the resulting loss of the important low frequency information. Furthermore, the same model order must match a smaller frequency band, giving a better fit.

Chapter 4

Accelerometer and Gyro Failure Detection

4.1 Introduction

In this chapter we discuss the sensor failure detection experiments conducted on some accelerometers and gyros of the Mini-Mast. These experiments are similar in nature to the experiments discussed in the previous chapter. Because we are using different types of sensors, we will get the interesting case where sensors of mixed type are used to form a double sensor parity relation. Three sensors are considered: two accelerometers that measure linear acceleration in the global X and Y directions, and the Z -axis gyro. All the sensors are at the tip of the mast. No state-space model was available for this set of sensors so we present only identified relations.

Before we discuss the failure detection experiments we first look at the spectra of the mea-

surements. In Figure 4.1 we show the spectrum of the Y -axis acceleration. The torque wheel




Figure 4.1: Spectrum of the Y -axis acceleration. The torque wheel motors were driven by discrete-time white noise that was passed through lowpass filters with 20 Hz bandwidth.

motors were driven by 20 Hz bandlimited random signals. We see that the first torsional mode is the dominant mode, with the first and second bending modes approximately 80 and 10 dB down respectively. We also see the effect of the local modes at 15 Hz and 19 Hz. The peaks in the spectrum at approximately 9 Hz and 23 Hz are probably the result of aliasing: the peak at 23 Hz is caused by the fourth bending mode at 43.4 Hz and the peak at 9 Hz comes from a mode at 74.8 Hz. Similarly, there are modes at 91.7 Hz and 93.2 Hz that alias to 25 Hz and 26.5 Hz respectively. In this experiment the sensor signals were filtered by third order analog lowpass filters before they were sampled, but the filtering was not enough to

prevent aliasing. The bandwidth of the analog filters was 20 Hz.

The solid line in Figure 4.2 shows the spectrum of the same measurement, but this time




Figure 4.2: Spectrum of the Y-axis acceleration. The torque wheel motors were driven by discrete-time random signals that were held constant for 4 sampling periods. The dashed line shows the spectrum when the sampling period is $2T_s$.

the torque wheel motors were driven by random signals that were held constant for 4 sampling periods. The output was sampled at $1T_s$ intervals, which corresponds to a sampling frequency of 66.67 Hz. Again we see the peaks at approximately 9 and 23 Hz. The dashed line in this figure shows the spectrum when we sample the output of the Y-axis accelerometer at $2T_s$ (33.33 Hz). Here we clearly see how the local mode at 19 Hz aliases to approximately 14 Hz. Although it was believed that the sensor outputs were filtered by 20 Hz analog filters before they were sampled, it was found after the experiments were conducted that the analog

filters were inadvertently set to have 100 Hz bandwidths, which is way above the sampling frequency. Although most of the aliased components are 30 dB or more down, it was found that the ambiguity caused by their presence degraded the performance of the Residual Generators. So all measurements were digitally filtered with a fifth order elliptical filter with 0.5 dB passband ripple and stopband attenuation of 40 dB; the equivalent continuous-time cutoff frequency was 7 Hz. The passband of this filter was chosen to be wide enough to pass the first five modes of the Mini-Mast and still give acceptable attenuation of the 9 Hz aliased component. The spectrum of the Z-axis gyro signal is shown in Figure 4.3 where




Figure 4.3: Spectrum of the Z-axis gyro signal. The torque wheel motors were driven by discrete-time random signals that were held constant for 4 sampling periods.

we see that the first torsional mode is by far the dominant mode.

4.2 Single Sensor Parity Relations

For this set of experiments the torque wheel motors were driven by random signals that were held constant for $4T_s$, while the sensor outputs were sampled at $1T_s$ intervals. A block diagram of the experimental setup is shown in Figure 4.4. In the block diagram we




Figure 4.4: Experimental setup.

also show an additional filter at the output of the residual generator. In some experiments we will show how additional filtering of the residuals can be used to improve the failure signature. This filter will be called the postfilter and we will indicate when it is used. A sixth order elliptical filter with 10 Hz bandwidth, 0.5 dB passband ripple and 60 dB stopband attenuation will be used in all the cases.

Figure 4.5 shows the failure of the Y-axis accelerometer at sample number 245 and Figure 4.6 shows a failure of the Z-axis gyro at sample number 255. In both cases identified SSPRs with 20 lags were used. Although both residuals indicate the corresponding failures, they contain high frequency noise and clearly will not give reliable indications of failures. The same residuals of Figures 4.5 and 4.6 are shown in Figures 4.7 and 4.8, but this time after the residuals were filtered by the postfilter. We see that lowpass filtering the residuals definitely leads to improved failure signatures. Figures 4.9 and 4.10 show the same sensors




Figure 4.5: SSPR for Y-axis accelerometer failure, 20 lags, $1T_s$. The Y-axis accelerometer has failed to zero at sample number 245.




Figure 4.6: SSPR for Z-axis gyro failure, 20 lags, $1T_s$. The gyro has failed to zero at sample number 255.




Figure 4.7: SSPR for Y-axis accelerometer failure, 20 lags, $1T_s$. This residual was filtered with the postfilter. The Y-axis accelerometer has failed to zero at sample number 245.




Figure 4.8: SSPR for Z-axis gyro failure, 20 lags, $1T_s$. The residual was filtered with the postfilter. The gyro has failed to zero at sample number 255.




Figure 4.9: SSPR for Y-axis accelerometer failure, 20 lags, $2T_s$. The Y-axis accelerometer has failed to zero at sample number 117.




Figure 4.10: SSPR for Z-axis gyro failure, 20 lags, $2T_s$. The gyro has failed to zero at sample number 137.

with the same type of failures at approximately the same points in time, but this time using a sampling period of $2T_s$. These two figures must be compared with Figures 4.5 and 4.6 respectively. First we note that the residuals have respectively 5 and 7.5 times larger amplitudes. Furthermore, the ratios of the amplitudes in the failed and unfailed states have increased considerably. The postfilter has not been applied to these residuals: the improvement comes only from the increased sampling period. It was found that filtering these residuals with the postfilter resulted in little improvement of the failure signature. In the next section we look at double sensor parity relations.

4.3 Double Sensor Parity Relations

Although the single sensor parity relations at $2T_s$ gave good performance there were failures where the indications were only marginal. The next step is to look at double sensor parity relations and hope that they will perform better. Figure 4.11 shows a failure of the X -axis accelerometer at sample number 236 and Figure 4.12 shows a failure of the Z -axis gyro at sample number 286. The number of lags used was (11,10), i.e., the parity relations had the form r_{ij} as shown in Equation (2.37), and we use the notation (i,j) to indicate the number of lags used. In both cases the unfiltered DSPR residuals are shown. We now have the interesting case where sensors of mixed type are used to construct the parity relations. The residual at the top of Figure 4.11 used the X -axis and Y -axis accelerometer measurements to compute the residual, while the residual shown at the bottom of this figure was computed from the X -axis accelerometer and Z -axis gyro measurements. A comparison of these residuals with their 20 lag, $1T_s$ single sensor counterparts (Figures 4.5 and 4.6) shows that we get a significant improvement by using the double sensor parity relations. It is again possible to clean up the signals with the postfilter but we will not show the results here.

An increase in the sampling period again leads to a significant improvement in the failure signatures as shown in Figures 4.13 and 4.14. Note that the output of the Residual Generators are shown in these figures: no extra filtering was applied to the residuals. In Figure 4.15 we have simulated the failure of an accelerometer that gives the correct output when the acceleration is positive and zero when the acceleration is negative. This type of failure can occur when a sensor is powered by a dual rail power supply and the negative supply falls away. The residuals clearly indicate this type of failure.

Despite the good results that we have shown so far, the Generalized Parity Relations are




Figure 4.11: DSPR for X-axis accelerometer failure, $(11,10)$ lags, $1T_s$.
 Top: residual with sensor pair (X-axis accelerometer, Y-axis accelerometer);
 Bottom: residual with sensor pair (X-axis accelerometer, Z-axis gyro).
 The X-axis accelerometer has failed to zero at sample number 236.





Figure 4.12: DSPR for Z-axis gyro failure, (11,10) lags, $1T_s$. The gyro has failed to zero at sample number 286. The residual at the top was constructed from the pair (X-axis accelerometer, Z-axis gyro) and the residual at the bottom from the pair (Y-axis accelerometer, Z-axis gyro).





Figure 4.13: DSPR for X-axis accelerometer failure, (11,10) lags, $2T_s$.
 Top: residual with sensor pair (X-axis accelerometer, Y-axis accelerometer);
 Bottom: residual with sensor pair (X-axis accelerometer, Z-axis gyro).
 The X-axis accelerometer has failed to zero at sample number 118.





Figure 4.14: DSPR for Z -axis gyro failure; (11,10) lags, $2T_s$. The gyro has failed to zero at sample number 143. The residual at the top was constructed from the pair (X-axis accelerometer, Z -axis gyro) and the residual at the bottom from the pair (Y-axis accelerometer, Z -axis gyro).





Figure 4.15: DSPR for X -axis accelerometer failure, $(11,10)$ lags, $2T_s$. The output of the X -axis accelerometer was half wave rectified from sample number 125 till the end.

Top: residual with sensor pair (X -axis accelerometer, Y -axis accelerometer);

Bottom: residual with sensor pair (X -axis accelerometer, Z -axis gyro).

still sensitive to certain parameter variations. In Figure 4.16 we show a failure of the Z -axis gyro at sample number 250. The torque wheel motors were driven by lowpass filtered white noise. The coefficients of the single sensor parity relation were identified on a different input-output data set, but with the motors driven by a similar type of input signal. Figure 4.17 shows a failure of the same sensor at the same time, using the same data set. However, the parity relation coefficients were identified on an input-output data set where the motors were driven by random signals that were held constant for $4T_s$. We see that the residual gives no indication of the failure. Repeating this test on the accelerometers gave the same result, i.e., no indication of failures. One possible explanation is that the torque wheel motors have a significant amount of friction so that the amplitudes of the input signals will determine how much the wheels are actually excited. The amplitudes of the lowpass filtered input signals were approximately 7 times smaller than the amplitudes of the input signals that were held constant. It is therefore difficult to conclude whether the difference in performance is due to the different type of input signals that were used or due to the different magnitudes of the input signals. Either case, it is a disturbing fact that the parity relations show this sensitivity to the different input signals.




Figure 4.16: SSPR for Z-axis gyro failure, 20 lags, $1T_s$. The gyro has failed to zero at sample number 250.




Figure 4.17: SSPR for Z-axis gyro failure, 20 lags, $1T_s$. The gyro has failed to zero at sample number 250.

4.4 Summary

In this chapter we have discussed the detection of accelerometer and gyro failures. It was found that the wider bandwidth of the measured signals can lead to aliasing that in turn degrades the performance of the residual generators. With proper anti-alias filters in place, the double sensor parity relations give good failure signatures. The sampling period again proved to be a very important parameter in the design of the Residual Generator. Despite the good performance, the parity relations are still very sensitive to the type and/or magnitude of the signals that are used to excite the system.

We also showed examples of parity relations that were constructed using different types of sensors. In all the cases considered the double sensor parity relations gave clear indications of all the different failure modes. It must be noted that this improved performance comes with the burden of an increased computational load.

It must be noted that we have shown results using parity relations with 20 lags throughout this chapter. It was found that, because this set of sensors have higher bandwidths than the displacement sensors, lower order models simply did not give clear indications of the failures.

Chapter 5

Actuator Failure Detection

In this chapter we discuss the detection of actuator failures on the Mini-Mast. For the experiments conducted here, the torque wheel motors were driven by lowpass filtered random signals. The bandwidths of these filters were 10Hz, and the baseline sampling period of 0.015 seconds, i.e., $1T_s$ was used. The measurements were filtered by 20Hz third order analog filters before they were sampled and digitized. We will present data only on results where the Bay 18 displacement sensors were used to obtain measurements, as the results obtained by using the accelerometers and gyro were similar in nature.

A failure was simulated while the experiment was conducted by disconnecting the commanded signal to a torque wheel motor. The model-based single actuator parity relation for this failure is shown in Figure 5.1. The actual time of failure is not known but should be approximately at sample number 500. In the figure we see that there is no indication of the failure at all. The residual of the same failure is shown in Figure 5.2 but this time an identified parity relation with 20 lags was used to detect the failure. Even though this residual is significantly smaller than the model-based residual, no indication of the failure




Figure 5.1: Model-based SAPR for X-wheel failure.





Figure 5.2: Identified SAPR for X-wheel failure.

is visible.

The above experiment was repeated by using double actuator parity relations, and both $1T_s$ and $2T_s$ sampling periods were used without any visible improvement. Single and double actuator parity relations were also identified using the X and Y -axes accelerometers and Z -axis gyro but they, too, were unable to detect the failure.

To gain more insight into the behavior of the actuator parity relations a computer simulation was conducted. Bandlimited random input signals were generated and a failure of the Y -torque wheel motor was simulated in the input data by zeroing the actual signal going to the plant. The Y -torque wheel motor was zeroed between samples number 213 and 284. The corresponding outputs were generated using the Mini-Mast state-space model given in Appendix B. The SAPR residual r_2 is shown in Figure 5.3.

This simulation was repeated, but this time noise was added to the measurements before they were used by the Residual Generator. A block diagram of this is shown below.



Actuator failure simulation.

The standard deviation of the noise that was added to the measurement was 1% of the standard deviation of the measurement. The SAPR residual for this simulation is shown in Figure 5.4. A comparison of the magnitudes of the residuals in Figures 5.3 and 5.4 shows




Figure 5.3: SAPR residual for Y-torque wheel failure. In this simulation the torque wheel was in a failed state between samples 213 and 284.

that the single actuator parity relations are extremely sensitive to noisy measurements. This sensitivity is also clearly visible when we look at the contributions of the measurements, r_y , and control inputs, r_u , to the residual r_2 as shown in Figure 5.5. In these figures we see that the noise in the measurement is amplified so much that it is orders of magnitude larger than the contribution of the control signal r_u . The extreme sensitivity to noise is easily explained when we look at the transfer functions of the corresponding Residual Generator, shown in Figure 5.6. In this figure we see that the transfer functions from the measurements y to the residual r_2 have very large gains over a large portion of the frequency band, especially at high frequencies, and therefore the smallest amount of noise in the measurements will be amplified and bury the residual deep in it. The figure also shows that the gain of the transfer function from the control signal u_2 to the residual r_2 is small compared to the gains




Figure 5.4: SAPR residual for Y -torque wheel failure with noisy measurements. The standard deviation of the added noise is 1% of the standard deviation of the measurement. The torque wheel was in a failed state between samples 213 and 284.

of the other transfer functions. Simulations with double actuator parity relations showed similar sensitivity to noise and gave no improvement.





Figure 5.5: Contribution r_y and r_u to the SAPR residual for Y-torque wheel. The standard deviation of the noise is 1% of the standard deviation of the measurement. The torque wheel was in a failed state between samples 213 and 284.




Figure 5.6: SAPR Residual Generator transfer functions for r_2 . B_2 is the transfer function from u_2 to r_2 , and A_{21} , A_{22} , A_{23} the transfer functions from y_1 , y_2 , y_3 to r_2 .

To summarize, detecting actuator failures on the Mini-Mast using Generalized Parity Relations was without any success. The main reason for the poor performance of the parity relations is the extreme sensitivity to noise, a result of the very high gains on the transfer functions from the measurements to the residual. Also, the small contribution of the control signal to the residual makes its absence very hard to detect. This sensitivity is inherent in the formulation of actuator parity relations.

Chapter 6

Generalized Parity Relation Experiments Using the ACES Facility

6.1 Introduction

The ACES (Active Control Evaluation for Spacecraft) facility is at the Marshall Space Flight Center in Huntsville, Alabama. It is a symmetrical beam, 13 m long, with a triangular cross section. Three longerons extend the full length of the mast and form the corners of the beam. There are 91 flexible batons in compression along the length of the boom connected by diagonal members in tension. The mast has a twist of 260 degrees from base to tip. An antenna is at the tip of the mast and a pointing gimbal at the base. A laser beam, fixed in the lab, is reflected by two mirrors and its $X - Y$ location is detected by an optical detector mounted on the antenna. One mirror is mounted on the antenna and the other one on the

pointing gimbal system or image motion compensator (IMC).




Figure 6.1: Schematic diagram of the ACES mast.

The mast is equipped with a variety of sensors and actuators and only the ones that are applicable to this work will be mentioned. Three-axes rate gyros and accelerometers are mounted at both the base and the tip of the mast. Two linear momentum exchange devices

(LMED) are located 6.4 and 11.4 m from the base respectively. The LMEDs consist of a proof mass actuator and collocated accelerometers. A high precision, three-axes, gimbal system with torque motors is mounted to the base of the system.

6.2 Processing the residuals

In the following section and the next chapter we discuss various failure detection experiments that require the comparison of the residual to a threshold. During the course of the work it was found that the failure could often be seen in the residual by the human eye but no significant increase in the magnitude was present so that it could be accurately detected by a threshold detector. Simple lowpass filtering turned out to be inadequate to extract the failure signature from the noise and it was found that simple nonlinear processing greatly improved the failure signature. A block diagram of the processing system is shown in Figure 6.2. The residual is first filtered with a lowpass filter with cut-off frequency of




Figure 6.2: Block diagram of post processing filter.

approximately 3 Hz. This cut-off frequency is based on the bandwidth of the system and should be large enough to include the dominant modes. In a few cases, notably ACES experiments, it was necessary to use a bandpass filter to remove some sensor biases. The output of this lowpass filter is then squared and applied to a second lowpass filter with

very low cut-off frequency, typically less than 0.05 Hz. This low frequency is necessary to ensure that the magnitude of the failure signature stays large during periods of small measurements. Butterworth and Bessel filters were initially used for this low bandwidth filter but numerical problems ruled them out. A simple moving average filter was found to work best and was used for all the experiments. The number of lags of this filter was typically 200 or more.

6.3 ACES Failure Experiments

The ACES failure detection experiments were conducted similar to the Mini-Mast experiments. In order not to excite too many high frequency modes bandlimited inputs were used on all the experiments. These input signals were generated by filtering random sequences with bandpass filters. Like the Mini-Mast, it was found that a longer sampling period gives better results. Because this was anticipated, the experiments were designed so that it would be easy to use various sampling periods by holding the input signal constant for more than one sampling period. Experiments were repeated by keeping the input signal constant for 1, 2, 4, and 8 sampling periods. The baseline sampling period for ACES is $T_s = 20$ ms. It was found that increasing the sampling period more than $4T_s$ lead to little improvement and the following results were all done at $4T_s$. The bandwidths of the bandpass filters were chosen so that the bandwidths of the resulting signals, taking the longer sampling period into account, were 1 Hz.

Initial results on ACES were very poor and it was hard to detect sensor failures using low order models. Only when models of order 60 and higher were used did failures show up in the residuals. Though failures were visible in the residuals by the eye, the residuals were not good enough to be used as reliable indicators of the failures. As explained in the previous

section, nonlinear filtering greatly improved the situation and was used in all the ACES experiments. The order of the parity relations presented here was 60.

All the parity relations used in these experiments were identified on one set of input-output data and then applied to other sets to detect component failures. Two different methods were used to identify the coefficients of the parity relations. The one approach was to use the standard least squares technique as outlined in Section 3.3. See also Section 2.1. For the second approach we used the ARX procedure in Matlab's System Identification Toolbox [9]. This procedure performs robust identification by weighing small errors quadratically and large errors linearly. The value at which the cost changes from quadratic to linear is user definable and the default value was used. Detail about this method can be found in [7] and [9]. Unless noted otherwise all parity relations were identified using the robust method.

For this facility there are many sensors and actuators to choose from and only a representative sample is showed here. The first SSPR was designed to detect failure of one of the gyros at the base of the mast. Recall that the mast is hanging upside down so the base is at the top and the tip at the bottom. The three base gimbal torque motors were used to excite the mast. A failure of the gyro was simulated by zeroing the recorded output from $t = 60$ seconds until the end of the experiment and the residual is shown in Figure 6.3. The coefficients for this residual were identified using the least squares technique. A sudden increase in the residual magnitude shows the failure and even after the transient has decayed a large residual still shows the presence of the failure.

SSPRs were also identified for the LMED accelerometers. A typical failure of one of the accelerometers is shown in Figure 6.4. The failure was introduced after 48 seconds. A clear indication of the failure is given by the increase in magnitude of the residual.

Due to the initial difficulties experienced in detecting sensor failures, single sensor parity




Figure 6.3: SSPR residual for base gyro. The failure was introduced at $t = 60$ seconds.




Figure 6.4: SSPR residual for LMED accelerometer. The failure was introduced at $t = 48$ seconds.

relations were abandoned in favor of double sensor parity relations with the hope that the additional measurement would give better results. With the nonlinear processing the SSPRs were found to compare favorably with the DSPRs and thus may be a viable solution for most sensors instead of the (numerically) more complex double sensor relations.

A double sensor parity relation was designed using a base and tip gyro with the base torquers as actuators. The tip gyro was failed to zero after 50 seconds; the residual is shown in Figure 6.5 where the increase in the magnitude of the residual clearly shows the failure. Failing the base gyro at $t = 40$ seconds also results in a clear indication of the failure as shown in Figure 6.6. Note that a set of three double sensor parity relations constructed from three different sensors is needed to determine which one of the sensors actually failed.

Combining different types of sensors in a DSPR is always interesting so a gyro at the base was combined with a tip accelerometer. A failure of the base gyro was introduced at $t = 60$ seconds and the residual is shown in Figure 6.7. The scale of the plot was kept the same as in Figure 6.6. Though the increase in magnitude is smaller than before the failure is still visible. Failing the accelerometer at $t = 45$ seconds is also detected as shown in Figure 6.8. Once the transient has decayed the magnitude of the residual is about the same as it was before the failure, which is not desirable. Still, the transient lasted long enough to be detected.

For the next DSPR we combined a tip gyro with an LMED accelerometer and failed the tip gyro after 50 seconds. The failure signature shown in Figure 6.9. The failure signature again gives a good initial indication of the failure yet it fails to keep the residual large till the end. A failure of the LMED accelerometer could not be detected. Analysis of the contributions of the individual sensors to the residual showed that the contribution from the accelerometer was more than 4 times smaller than that of the gyro and was effectively




Figure 6.5: DSPR residual for (base gyro, tip gyro) pair. The tip gyro failure was introduced at $t = 50$ seconds.




Figure 6.6: DSPR residual for (base gyro, tip gyro) pair. The base gyro failed at $t = 40$ seconds.




Figure 6.7: DSPR residual for (base gyro, tip accelerometer) pair. The base gyro failed at $t = 60$ seconds.




Figure 6.8: DSPR residual for (base gyro, tip accelerometer) pair. The tip accelerometer failed at $t = 40$ seconds.




Figure 6.9: DSPR residual for (tip gyro, LMED accelerometer) pair. The tip gyro failed at after 50 seconds.

masked by the noise. The transient seen in Figure 6.9 resulted from the transient that occurred when the gyro was failed. Once this is zero the contribution from the other sensor is so small that there is no sustained indication of the gyro failure. This also explains why the gyro failure in Figure 6.7 resulted in a transient with no sustained indication of the failure.

It was mentioned that different identification procedures were used to identify the coefficients of the parity relations. We will now show how the choice of method affects the quality of the residual. A DSPR was identified for the base gyro and tip accelerometer pair that was showed in Figure 6.8. This time the standard least squares method was used and the resulting failure is showed in Figure 6.10 where we see that the failure is not detected




Figure 6.10: DSPR residual for (base gyro, tip accelerometer) pair. The accelerometer failed at after 40 seconds.

by this DSPR. This behavior was not general and in some cases DSPR identified with the standard least squares method resulted in better residuals!

Similar results were obtained when the LMEDs were used to excite the mast instead of the base torquers. Experiments were also conducted where only the IMCs were used with the hope that the direct link between the IMCs and detector will give good results. The large quantization error of the detector made this the worst performer of the sensor-actuator pairs tested. Combining the detector with any set of actuators always gave bad results and detector failures could not be detected.

Both single and double actuator parity relations were identified from input-output data but actuator failures could not be detected. In all the cases the individual contributions

to the residual from the sensors were much larger than the contribution from the actuator. Typical results are shown in Figure 6.11. An SAPR was constructed for the X -gimbal torquer using two tip gyros as sensors. The individual contributions are shown Figure 6.11. In this experiment the X -gimbal torquer was disconnected between 60 and 100 seconds.

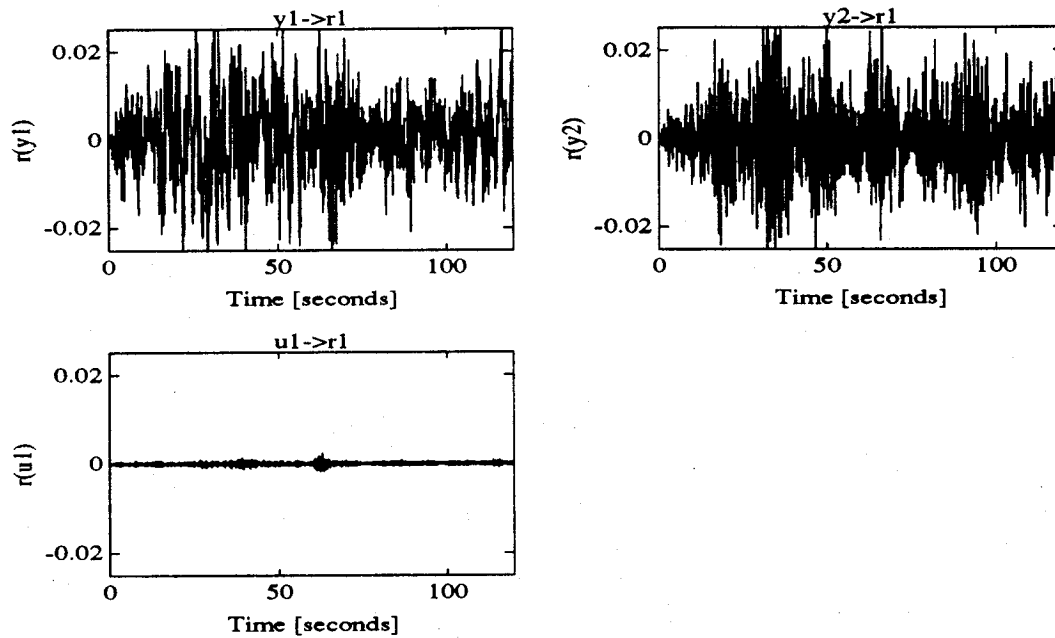


Figure 6.11: SAPR individual contributions for X -gimbal torquer using two tip gyros as sensors. The torquer was disconnected between 60 and 100 seconds.

No indication of this failure is given by these signals. As we showed in Chapter 5 (see for example Figure 5.6), the poor performance is a result of the extreme sensitivity of SAPR Residual Generators to noise.

6.4 Summary

Detecting sensor failures on the ACES facility proved to be difficult compared to the Mini-Mast. The main reason for this is the high order models that are needed to construct reliable parity relations. When models of order 60 and higher are used it is possible to detect sensor failures with both single and double sensor parity relations. In all the cases nonlinear processing of the residuals was needed to give reliable failure signatures.

For some double sensor parity relations the failure of one sensor would result in a transitory indication only with the contribution of the second sensor effectively masked by the noise. Currently no method exists to determine analytically which sensor pairs will give "good" residuals. It was also shown that robust identification techniques can identify parity relations that perform better than ones identified using standard least squares techniques.

Actuator failures proved to be hard to detect for this system and even high order models did not detect failures. The main reason for this is the extreme sensitivity of actuator parity relations to noise and the small contribution of the nominal actuator signal to the parity relation residual.

Chapter 7

Failure Detection Filters

7.1 Introduction

As stated in Chapter 1 there is a variety of methods available for the detection and identification of control system component failures. In previous chapters we looked at generalized parity relations; in this chapter we discuss a series of failure detection experiments that were conducted on the Mini-Mast using the failure detection filter (FDF). First, we give a brief description of failure detection filters.

The FDF is an observer of a linear dynamic system which indicates failures of some of the components of the system by constraining the signature of the failures to unique directions in the space of the measurement residuals. It was developed by Beard in 1971 [2]. Shortly after that Jones [6] gave a geometric interpretation of the failure detection filter and more recently White and Speyer [16] viewed it as an eigensystem assignment problem. A different formulation of the failure detection filter was developed by Massoumnia [10]. We will give

brief descriptions of the Beard and Massoumnia formulations and refer the reader to the references for more detail.

Consider the continuous-time, linear, time-invariant model given by

$$\dot{\mathbf{x}}_p(t) = \mathbf{A}_p \mathbf{x}_p(t) + \mathbf{B}_p \mathbf{u}_a(t), \quad (7.1)$$

$$\mathbf{y}(t) = \mathbf{C}_p \mathbf{x}_p(t) \quad (7.2)$$

where $\mathbf{x}_p(t) \in \mathbb{R}^{n_x}$ is the state vector, $\mathbf{u}_a(t) \in \mathbb{R}^{n_u}$ is the actual input vector driving the system, $\mathbf{y}(t) \in \mathbb{R}^{n_y}$ is the measurement vector, and $\mathbf{A}_p \in \mathbb{R}^{n_x \times n_x}$, $\mathbf{B}_p \in \mathbb{R}^{n_x \times n_u}$, and $\mathbf{C}_p \in \mathbb{R}^{n_y \times n_x}$ are the usual continuous-time state-space matrices. The failure detection filter is a linear time-invariant system driven by the commanded inputs \mathbf{u}_c and measured outputs and is described by

$$\dot{\mathbf{z}}(t) = \mathbf{A}_f \mathbf{z}(t) + \mathbf{D} \mathbf{y}(t) + \mathbf{B}_f \mathbf{u}_c(t), \quad (7.3)$$

and the matrices $\mathbf{A}_f \in \mathbb{R}^{n_x \times n_x}$, $\mathbf{B}_f \in \mathbb{R}^{n_x \times n_u}$, and the detector gain $\mathbf{D} \in \mathbb{R}^{n_x \times n_y}$ must be selected to produce the necessary information about the failure. This can be done by constructing the error signal

$$\mathbf{e}(t) = \mathbf{x}_p(t) - \mathbf{z}(t), \quad (7.4)$$

which is just the difference between the system states and the filter states. Differentiating this equation and using Equations (7.1)-(7.3) we get

$$\dot{\mathbf{e}}(t) = (\mathbf{A}_p - \mathbf{D} \mathbf{C}_p) \mathbf{x}_p(t) - \mathbf{A}_f \mathbf{z}(t) + (\mathbf{B}_p \mathbf{u}_a(t) - \mathbf{B}_f \mathbf{u}_c(t)). \quad (7.5)$$

If the actual control signal driving the plant is the same as the commanded control signal we see that by choosing

$$\mathbf{A}_f = \mathbf{A}_p - \mathbf{D} \mathbf{C}_p \quad (7.6)$$

$$\mathbf{B}_f = \mathbf{B}_p, \quad (7.7)$$

we get the following expression for the error dynamics

$$\dot{e}(t) = (A_p - DC_p)e(t). \quad (7.8)$$

We also define the residual

$$r = C_p e(t). \quad (7.9)$$

If we choose D so that all the poles of $A_p - DC_p$ are in the left half plane then the residual will go to zero as time goes to infinity, that is, if everything in the system is in good health the steady-state residual will be zero.

Consider now a failure of the j th actuator modelled by

$$u_a(t) = u_c(t) + i_j n(t) \quad (7.10)$$

where i_j is the j th column of the identity matrix and $n(t)$ is an arbitrary unknown function of time. Substituting Equation (7.10) into Equation (7.5) and using the choices for A_f and B_f given in Equations (7.6) and (7.7) we get the following expression for the error dynamics

$$\dot{e}(t) = (A_p - DC_p)e(t) + b_j n(t) \quad (7.11)$$

where b_j is the j th column of B_p . From this equation we see the error is now driven by the signal $n(t)$ so that the error, and therefore the residual, is nonzero in general. Beard has shown that it is possible to find a detector gain D so that the steady-state residual maintains a fixed direction, determined by C_p and b_j , in the output space and the eigenvalues of $A_p - DC_p$ can be assigned almost arbitrarily. There are also other important concepts related to the design of the filter, e.g., it is possible to design a filter that will detect the failure of more than one actuator in which case it is important to determine whether a group of failures is output separable and mutually detectable. It is also possible to design a filter that will detect sensor failures in which case the residual is confined to a plane in the output space. We will not discuss the detail here and the reader is referred to Beard [2]. The

important point is that it is possible to detect the presence of a failed component as well as identify which one actually failed. This can be done, in principle, for arbitrary functions $n(t)$ — which means arbitrary modes of failure.

A variation of the failure detection filter was proposed by Massoumnia [10]. To illustrate the method, assume we have a system with two inputs and we would like to design a filter that is sensitive to a failure of the first actuator but not of the second. The system and failure are modelled by

$$\dot{x}_p(t) = A_p x_p(t) + B_p u(t) + l_1 m_1(t) + l_2 m_2(t) \quad (7.12)$$

$$y(t) = C_p x_p(t) \quad (7.13)$$

and the filter and residual are given by

$$\dot{w}(t) = A_f w(t) - B_{fy} y(t) + B_{fu} u(t) \quad (7.14)$$

$$r(t) = C_f w(t) - D_{fy} y(t) \quad (7.15)$$

where $A_f \in \mathbb{R}^{n_x \times n_x}$, $B_{fy} \in \mathbb{R}^{n_x \times n_y}$, $B_{fu} \in \mathbb{R}^{n_x \times n_u}$, $C_f \in \mathbb{R}^{n_r \times n_x}$, and $D_{fy} \in \mathbb{R}^{n_r \times n_y}$, and n_r is the dimension of the residual. In this model the term $l_1 m_1(t)$ represents the component whose failure we would like to detect with this filter. This term is analogous to the way we simulated the actuator failure for the Beard filter — when the actuator is fully operational $m_1(t)$ is zero and when a failure is present $m_1(t)$ is an arbitrary function of time. The term $l_2 m_2(t)$ models the failure of the other actuator that should not influence the residual of this filter. The functions $m_1(t)$ and $m_2(t)$ are called failure modes. The idea is to find the filter parameters A_f , B_{fy} , B_{fu} , C_f , and D_{fy} so that the residual is nonzero when $m_1(t)$ is nonzero and zero when $m_1(t)$ is zero, independent of $m_2(t)$. Expressed in a different way, we would like the transfer function from $m_2(t)$ to $r(t)$ to be zero and $m_1(t)$ should be input observable. The Massoumnia filter is more general than we have presented it here and can

handle systems with any number of inputs. Also, the vectors \mathbf{l}_1 and \mathbf{l}_2 in Equation (7.12) can in general be matrices \mathbf{L}_1 and \mathbf{L}_2 so that it is possible to design a filter that will detect a set of failures while not being influenced by another set of failures. The general description for the system thus is

$$\dot{\mathbf{x}}_p(t) = \mathbf{A}_p \mathbf{x}_p(t) + \mathbf{B}_p \mathbf{u}(t) + \mathbf{L}_1 \mathbf{m}_1(t) + \mathbf{L}_2 \mathbf{m}_2(t). \quad (7.16)$$

There is also some freedom in the selection of the filter parameters and Massoumnia has shown how this can be used to whiten the residuals in the absence of failures. The conditions under which this problem can be solved can be found in [10].

The failure detection filter formulated by Massoumnia can also be used to detect sensor failures. Here we model the plant and sensor failures by

$$\dot{\mathbf{x}}_p(t) = \mathbf{A}_p \mathbf{x}_p(t) + \mathbf{B}_p \mathbf{u}(t) \quad (7.17)$$

$$\mathbf{y}(t) = \mathbf{C}_p \mathbf{x}_p(t) + \mathbf{J}_1 \mathbf{n}_1(t) + \mathbf{J}_2 \mathbf{n}_2(t) \quad (7.18)$$

where we would like to detect sensor failures modelled by the term $\mathbf{J}_1 \mathbf{n}_1(t)$ but not those modelled by the term $\mathbf{J}_2 \mathbf{n}_2(t)$. To design a filter that will detect a failure of the j th sensor but not respond to failures in any of the other sensors we choose $\mathbf{J}_1 = \mathbf{i}_j$ where \mathbf{i}_j is the j th column of the identity matrix and \mathbf{J}_2 is all but the j th column of the identity matrix. We will model the signals $\mathbf{n}_1(t)$ and $\mathbf{n}_2(t)$ by the following systems

$$\dot{\mathbf{x}}_{n_1}(t) = \mathbf{A}_1 \mathbf{x}_{n_1}(t) + \mathbf{B}_1 \mathbf{m}_1(t) \quad (7.19)$$

$$\mathbf{n}_1(t) = \mathbf{C}_1 \mathbf{x}_{n_1}(t) \quad (7.20)$$

$$\dot{\mathbf{x}}_{n_2}(t) = \mathbf{A}_2 \mathbf{x}_{n_2}(t) + \mathbf{B}_2 \mathbf{m}_2(t) \quad (7.21)$$

$$\mathbf{n}_2(t) = \mathbf{C}_2 \mathbf{x}_{n_2}(t). \quad (7.22)$$

We can now form the augmented system

$$\begin{pmatrix} \dot{x}_p(t) \\ \dot{x}_{n_1}(t) \\ \dot{x}_{n_2}(t) \end{pmatrix} = \begin{pmatrix} A_p & 0 & 0 \\ 0 & A_1 & 0 \\ 0 & 0 & A_2 \end{pmatrix} \begin{pmatrix} x_p(t) \\ x_{n_1}(t) \\ x_{n_2}(t) \end{pmatrix} + \begin{pmatrix} B_p \\ 0 \\ 0 \end{pmatrix} u(t) + \begin{pmatrix} 0 \\ B_1 \\ 0 \end{pmatrix} m_1(t) + \begin{pmatrix} 0 \\ 0 \\ B_2 \end{pmatrix} m_2(t) \quad (7.23)$$

$$\triangleq Ax(t) + Bu(t) + L_1 m_1(t) + L_2 m_2(t) \quad (7.24)$$

$$y(t) = [C_p \quad J_1 C_1 \quad J_2 C_2] x(t), \quad (7.25)$$

which is in the form of Equation (7.16). So we can use the same techniques as before to solve the sensor failure detection problem.

7.2 Actuator Failure Experiments

In this section we discuss failure detection experiments on the Mini-Mast where the failure detection filter was used. The Massoumnia filter was used in all the experiments. The main reason for this choice was the availability of numerically reliable software to compute the filter parameters. Although the Beard filter has the advantage that the filter can be designed to detect and isolate more than one failure the algorithm given to compute the detector gain D in Equation (7.3) gives numerical problems when applied to high order systems.

A twelfth order state-space model was identified from the recorded input-output data using

the method of [15]. In all the experiments the baseline sampling frequency was used so that the sampling period was $T_s = 15$ ms. Recall that for parity relations this short sampling period did not give good results for sensor failures, and no actuator failures could be detected. The sensors used were the three displacement sensors at the top of the mast, and the actuators were the three torque wheels mounted at the top of the mast.

Three failure detection filters were designed with whitening of the residuals. As an example for the first actuator, the system was modelled by

$$\dot{x}_p(t) = A_p x_p(t) + B_p u(t) + b_1 m_1(t) + [b_2, b_3] m_2(t), \quad (7.26)$$

where $B_p = [b_1, b_2, b_3]$ so that the corresponding residual r_1 will detect failures of the first actuator only. The second filter was designed to detect failures of the second actuator only and the system was modelled by

$$\dot{x}_p(t) = A_p x_p(t) + B_p u(t) + b_2 m_1(t) + [b_1, b_3] m_2(t) \quad (7.27)$$

with corresponding residual r_2 . The third actuator is treated in a similar way. The recorded input-output data was then used as inputs to the failure detection filters. Unless stated otherwise, the residuals were processed as shown in Figure 7.1 and outlined in Section 6.2.




Figure 7.1: Block diagram of post processing filter.

In Figure 7.2 we show the residual when the X-torque wheel was disconnected at approxi-




Figure 7.2: FDF residual r_{f1} for X -torque wheel failure. The failure existed approximately between $t = 30$ and $t = 60$ seconds.




Figure 7.3: FDF residual r_{p1} for X -torque wheel failure. The residual was processed as shown in Figure 6.2

mately 30 seconds and reconnected approximately 30 seconds later. The residual shown was taken off just after the first lowpass filter and before the nonlinear element, see Figure 6.2. An increase in the magnitude of the residual clearly shows the failure of the actuator. The residuals for the Y- and Z-torque wheels were not affected by this failure. Figure 7.3 shows the residual at the output of the moving average filter and it is clear from the large magnitude of this signal that we can detect the failure with a threshold detector. The residual for a failure of the Y-torque wheel is shown in Figure 7.4 where the residual was taken at the output of the moving average filter; the increase in the magnitude clearly shows the malfunctioning of the actuator. Again, the other residuals were not affected by the failure. The Z-torque wheel was also disconnected in another experiment and we show the residual, taken directly from the output of the FDF, in Figure 7.5 where again we have a good indication of the failure. Note that no additional filtering or processing was performed on this residual and lowpass filtering would be enough to give an excellent indication of this failure.

Because these actuator failures were not simulated by corrupting the recorded input-output data they are significant indicators that actuator failures *can* be detected for large flexible structures.

In Chapter 5 we mentioned that actuator parity relations fared poorly on the Mini-Mast. Since those experiments were performed we started using the post processing filter that greatly improved the ability to extract the failure signature from the noise. To make a fair comparison between the FDF and SAPR, a 12-lag SAPR has been identified and applied to the same input-output data; the residual was processed in the same way as we did for the FDF. Figure 7.6 shows the resulting failure signature — this figure should be compared with Figure 7.4. The SAPR simply fails to give any indication of the failure. One reason for the big difference in performance of the two methods is found by comparing the transfer




Figure 7.4: FDF residual r_{p2} for Y-torque wheel failure. The failure existed approximately between $t = 30$ and $t = 60$ seconds.




Figure 7.5: FDF residual r_3 for Z-torque wheel failure. The failure existed approximately between $t = 30$ and $t = 60$ seconds.




Figure 7.6: SAPR residual r_{p2} for Y -torque wheel failure. The failure existed approximately between $t = 30$ and $t = 60$ seconds.

functions of the two residual generators. Figures 7.7 and 7.8 show the transfer functions from the measurements $y_1 - y_3$ and input u_2 to the residual r_2 of the SAPR and FDF residual generators respectively. From the first of these figures we see the gains on the measurements of the SAPR residual generator increase with increasing frequency and thus will amplify high frequency noises that are present in the measurements. The corresponding gains of the FDF increase only slightly over the same frequency range and stay more or less constant at high frequencies. In Figure 7.8 we also show the transfer function from the failure mode $m_1(t)$ to the residual $r_2(t)$ and see that the FDF puts a little more emphasis on frequencies at 0.9 Hz, which corresponds to the first bending mode. In the next section we present the results on sensor failures using the FDF.




Figure 7.7: Residual generator transfer functions for Y-torque wheel SAPR.




Figure 7.8: Residual generator transfer functions for Y-torque wheel FDF.

7.3 Sensor Failure Experiments

Three FDFs were designed to detect failures of the three displacement sensors located at the tip of the Mini-Mast. Sensor failures were simulated in the recorded input-output data by corrupting the measurements. In all the experiments presented in this section the baseline sampling period of 15 ms was used and the sensors were failed to zero at $t=15$ seconds. Figure 7.9 shows the residual directly at the output of the detection filter, i.e., with no




Figure 7.9: FDF residual r_1 for Sensor D1 failure.

additional processing, when Sensor D1 has failed. The failure mode $m_1(t)$ was modelled by a first order system with a pole at $s = -2\pi 5$ rad/s. The processed residual also gave a clear indication of the failure and is not shown. None of the other residuals were affected by the failure of this sensor. Failures simulated for the other sensors also showed up clearly in the unprocessed residual and the results are not shown.

A second set of detection filters were designed to detect sensor failures but this time the dynamics of the failure modes were each modelled by second order systems with natural frequencies of $\omega_n = 4.4$ Hz and damping coefficients of $\zeta = 0.2$. The natural frequency, which was chosen to be approximately 10% faster than the frequency of the first torsional mode, was selected based on the strong dominance of this frequency in all the measurements. The resulting residual for a failure of Sensor D1 is shown in Figure 7.10. When we compare this




Figure 7.10: FDF residual r_1 for Sensor D1 failure. A second order model was used for the failure modes.

figure with Figure 7.9 we see that the slightly more complex model used for the dynamics of the failure mode resulted in a residual that gives a better indication of the failure. Note that no additional processing was done: this is the residual as computed by the FDF.

Motivated by the improvement in results obtained by modelling the dynamics of the failure mode with a second order system, the same idea was applied to the actuator failure detection

filter. That is, the failure mode was explicitly modelled by augmenting the model of the plant with the dynamics of the failure mode. Using the same values for the natural frequency and damping coefficient we repeated some of the experiments; the residual corresponding to a failure of the Y-torque wheel motor is shown in Figure 7.11. Figure 7.12 shows the unprocessed residual corresponding to Figure 7.4. This residual was generated with the standard actuator FDF, i.e., the failure mode was not modelled explicitly. Comparing Figures 7.11 and 7.12 we see that modelling the failure mode led to a residual that gives a better indication of the failure and only simple lowpass filtering is needed to extract the failure signature. The added complexity thus appears to be advantageous and further work is necessary to determine how useful this approach is. Also, it is necessary to determine whether this modelling will impair the filter's ability to detect any failure mode.

ACES

Several attempts were made to design failure detection filters for ACES. The FDF requires a state-space description of the system and as we have seen for ACES high order models are needed to detect component failures. Identifying such a model is not a trivial task and was beyond the scope of this work. A twelfth order model was tried without success.

7.4 Computational Burden

In this section we give a comparison of the approximate number of operations that is necessary to implement generalized parity relations and failure detection filters. The computation of

$$\sum_{i=1}^n a_i x_i \quad (7.28)$$

requires n multiplications and $n - 1$ additions. However, because the time to do an addition is usually much shorter than the time to do a multiplication, we will assume that n mul-




Figure 7.11: Unprocessed FDF residual r_2 for Y -torque wheel failure. The failure mode was modelled by a second order system.




Figure 7.12: Unprocessed FDF residual r_2 for Y -torque wheel failure.

tifications and additions are needed to compute the sum. Thus, calling a multiplication and addition an operation, n operations are needed to compute the given sum. Addition of two scalar quantities is ignored. Sums like the one given in Equation (7.28) occur in difference equations, e.g., parity relations, as well as in matrix-vector multiplications, e.g., the detection filter.

We consider residual generators designed to detect the failure of a single component. A general description for a single sensor parity relation using n_x lags is

$$r(t) = y(t) + \sum_{j=1}^{n_u} \sum_{s=1}^{n_x} \beta_{j,s} u_j(t-s) + \sum_{s=1}^{n_x} \alpha_s y(t-s) \quad (7.29)$$

and the number of operations necessary to implement this is approximately

$$n_{\text{sspr}} = n_x(n_u + 1). \quad (7.30)$$

For a single actuator parity relation the roles of sensors and actuators are interchanged so that

$$n_{\text{sapr}} = n_x(n_y + 1). \quad (7.31)$$

The general description for the failure detection filter is given by Equations (7.14) and (7.15), and straight forward implementation of this requires

$$n_{\text{fdf}} = n(n + n_y + n_u + 1) + n_y \quad (7.32)$$

operations for a one dimensional residual, where n is the dimension of \mathbf{A}_f . The filter is often of dimension less than n_x because the part of the dynamics that is unobservable from the residual is automatically factored out during the design of the filter. For a filter designed to detect sensor failures, the dynamics of the failure modes that must be ignored by the filter also gets factored out. In many cases complete columns of \mathbf{B}_{fy} , \mathbf{D}_{fy} , and \mathbf{B}_{fu} are zero because the failure detection filter ignores certain measurements and inputs. As an example,

for the Mini-Mast sensor filters, only the columns of B_{fy} and D_{fy} that correspond to the specific sensor the filter was designed for had nonzero entries. Furthermore, the number of operations can be decreased by diagonalizing the resulting A -matrix; if this is done and assuming all poles are complex, only

$$n_{\text{fdf-diag}} = n(n_y + n_u + 3) + n_y \quad (7.33)$$

operations are needed. It is also possible to implement the filter as a difference equation similar to that of a parity relation; in this case the factor 3 in the last expression changes to 1. The number of operations for the failure detection filter is more than that of the parity relations by approximately $n(n_y + 3) + n_y$, where n is of the same order as the dimension of the plant. The factor nn_y comes from computing the term $B_{fy}y$ and, as we noted earlier, often only one column of B_{fy} is nonzero so that the filter has approximately $3n$ more operations than the parity relation. This is a small price to pay for the increase in performance.

7.5 Summary

In this chapter we conducted several failure detection experiments using the failure detection filter. With some additional processing of the residual we showed that it is possible to get good indications of actuator failures on the Mini-Mast. Comparison of the transfer functions of the two residual generators also showed that the failure detection filter is less sensitive to noise. The actuator failures that were detected can be regarded as a good measure of the failure detection filter's ability to detect actuator failures as these were not simulated failures. Compared to single actuator parity relations the failure detection filter is the clear winner.

In the case of sensor failures the failure detection filter again performed better than single sensor parity relations. In previous chapters we have shown several times that generalized parity relations prefer a longer sampling period and often performed poorly at the baseline sampling period of the Mini-Mast. The failure detection filter was tested only at the baseline sampling period and it performed well for both actuator and sensor failures. This can be of great importance in applications where the control system requires fast sampling.

Chapter 8

Conclusion

Space based stations put strict requirements on the reliability of the control system components. Because these systems will be used for long unattended periods of time the control system must be tolerant to the failure of its actuators and sensors. The reliability of the system can be increased through hardware redundancy, but this leads to increased weight and can be impractical when many components are used by the control system. The reliability of the system can also be increased with analytical redundancy that uses the redundancy that is present in the dynamics of the plant and the input-output histories.

Ideally one would require that a failure detection and isolation system be independent of the mode of failure and it should also be applicable to both sensors and actuators. Two methodologies satisfy these requirements: the Failure Detection Filter and the method of Generalized Parity Relations. In this work we discussed the application of these two methods experimental flexible space structures, the NASA Langley Mini-Mast and the Marshall Space Flight Center ACES mast. Different sensor sets were considered and the detection of actuator failures was also investigated.

The performance of a reduced order model for the Mini-Mast that included the first five modes of the plant (referred to as model-based relations) was compared to a set of parity relations that was identified directly from input-output data (referred to as the identified relations). The effect of the model order and sampling period on the performance of the Residual Generator were also shown.

For the Mini-Mast, The first set of sensors consisted of the three displacement sensors at the tip of the mast. These sensors measured the displacement of the tip relative to a fixed, rigid structure that was built around it. The model-based residuals suffered from sensitivity to noise and did not give reliable indications of the failures. The identified relations gave good failure signatures on all the different failure modes that were simulated in the data. Because all analytical redundancy techniques use a model of the plant, they all suffer from mismatches between the model and the real plant. By identifying the coefficients of the parity relations directly from input-output data the need for an accurate state-space model of the plant disappears. Identifying the parity relations has the advantage that it is easy to increase the model of the order if the low order models that are typically used by the control system give unacceptable performance. Using double sensor parity relations led to no improvement for the model-based relations, while the identified relations showed a significant improvement in the failure signature. It was also illustrated that the sampling period had a significant effect on the performance of the Residual Generators; it was found that the longer sampling periods gave better failure signatures. The reason for this improved performance comes from the smaller frequency band that needs to be matched by a model with a given order.

The second sensor set for the Mini-Mast consisted of two accelerometers and a gyro, all mounted at the tip of the mast. A state-space model was not available for this set of sensors so all the results apply to identified relations. Because of the wider bandwidths

of these sensors it was found that we had to increase the order of the model to get good performance from the Residual Generators. The single sensor parity relations performed satisfactorily and the double sensor parity relations gave good failure signatures. Again, increasing the sampling period resulted in a significant improvement of the failure signatures. This combination of sensors also illustrated that it is possible to use sensors of mixed type to construct parity relations.

A set of parity relations that was identified when the Mini-Mast was excited by bandlimited signals performed poorly when applied to data that was recorded when the mast was driven by wideband signals. The magnitudes of the input signals differed considerably so that the poor performance is probably caused by the nonlinearities of the torque wheel actuators.

The detection of Mini-Mast actuator failures using Generalized Parity Relations proved to be very difficult. It was found that the Residual Generators had very large gains associated with the transfer functions from the measurements to the residual, making it extremely sensitive to noise in the measurements. This resulted in residuals that were so noisy that it completely obscured the contribution of the control input.

Identified Generalized Parity Relations were also used to monitor component failures on the ACES mast. Both single sensor and double sensor parity relations needed 60 or more lags to give acceptable performance. Nonlinear processing of the residuals improved the failure signatures significantly. In some of the double sensor parity relation experiments the failure of only one of the sensors could be detected because the contribution to the residual of the other sensor is very small and thus masked by the noise. We have also showed that the method used to identify the coefficients of the parity relation can have a significant effect on the performance of the parity relation.

The failure detection filter performed very well for both sensor and actuator failures on

the Mini-Mast. Some of the residuals for sensor failures were so good that no additional processing was necessary. While generalized parity relations failed to detect actuator failures on the Mini-Mast the failure detection filter detected all the failures. Furthermore, it was found that the failure detection filter worked very well with short sampling periods.

The work concluded with an approximate analysis of the computational burden for the two residual generators. The number of operations for the failure detection filter is more than that of the generalized parity relation by a linear term.

Future Work

Though FDI has been studied for many years, several problems remain unsolved. A brief summary of some problems that need further investigation is given here. It was pointed out that an increase in the sampling period led to improved failure signatures. Although no examples were given it was found that increasing the sampling period beyond a certain point yielded no improvement. It thus appears that there might be an optimum sampling period. Even if we can find such an optimal sampling period analytically, it may not be an acceptable sampling period for use by the controller. It is easy to derive parity relations when the sampling period used by the Residual Generator is an integer multiple of the sampling period used by the controller. However, the analysis of the system is complicated because the resulting Residual Generator is not time invariant any more. Because of the large improvement that can be realized by the selection of a good sampling period it is an area that warrants further investigation.

It was pointed out at the end of Section 2.2 that the construction of double sensor parity relations leads to a choice of two relations. An example of this was given in Section 3.6 where we saw that the use of the second relation gave better failure signatures. We have also showed that in some cases a double sensor parity relation fails to indicate the failure

of one of the sensors it was designed for. Analysis of double sensor parity relations may lead to additional insight to why this happened and maybe hint at which relations should be used for best results.

When doing model validation the residual is studied in great detail as this signal contains a wealth of information about the identified model. Thus we see that model validation is similar in nature to failure detection. Robust identification techniques are constantly being developed. Because analytical redundancy methodologies all rely on a model of the plant robustness is always an issue. The applicability of these robust identification techniques to failure detection must therefore be investigated.

Appendix A

Second order system analysis

In this appendix we will find the discrete-time description of a continuous-time second order system. Let the continuous-time system be

$$y(s) = \frac{\omega_n^2}{s^2 + 2\zeta\omega_n s + \omega_n^2} u(s), \quad (\text{A.1})$$

where ω_n is the natural frequency and $\zeta < 1$ the damping ratio. A continuous-time state-space description is (see Chen [3], chapter 6)

$$\dot{\mathbf{x}}(t) = \begin{pmatrix} 0 & -\omega_n^2 \\ 1 & -2\zeta\omega_n \end{pmatrix} \mathbf{x}(t) + \begin{pmatrix} \omega_n^2 \\ 0 \end{pmatrix} u(t), \quad (\text{A.2})$$

$$y(t) = [0 \quad 1] \mathbf{x}(t) \quad (\text{A.3})$$

$$= \mathbf{c}' \mathbf{x}(t). \quad (\text{A.4})$$

The discrete-time state-space description is given by Equations (2.8) – (2.11). Let the damped natural frequency be denoted by

$$\omega_d = \omega_n \sqrt{1 - \zeta^2}. \quad (\text{A.5})$$

Evaluating the equations for \mathbf{A} and \mathbf{B} we get

$$\mathbf{A} = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} \quad (\text{A.6})$$

$$= e^{-\zeta\omega_n T_s} \begin{pmatrix} \cos(\omega_d T_s) + \frac{\zeta\omega_n}{\omega_d} \sin(\omega_d T_s) & \frac{-\omega_n^2}{\omega_d} \sin(\omega_d T_s) \\ \frac{1}{\omega_d} \sin(\omega_d T_s) & \cos(\omega_d T_s) - \frac{\zeta\omega_n}{\omega_d} \sin(\omega_d T_s) \end{pmatrix}, \quad (\text{A.7})$$

and

$$\mathbf{B} = \begin{pmatrix} b_{11} \\ b_{21} \end{pmatrix} \quad (\text{A.8})$$

$$= \begin{pmatrix} 2\zeta\omega_n - e^{-\zeta\omega_n T_s} \left(2\zeta\omega_n \cos(\omega_d T_s) + \frac{\omega_n(2\zeta^2 - 1) \sin(\omega_d T_s)}{\sqrt{1 - \zeta^2}} \right) \\ 1 - e^{-\zeta\omega_n T_s} \left(\cos(\omega_d T_s) + \frac{\zeta \sin(\omega_d T_s)}{\sqrt{1 - \zeta^2}} \right) \end{pmatrix}. \quad (\text{A.9})$$

The numerator and denominator polynomials are found by evaluating Equation (2.68)

$$y(z) = \mathbf{c}'(z\mathbf{I} - \mathbf{A})^{-1} \mathbf{b}u(z) \quad (\text{A.10})$$

$$= \frac{n(z)}{d(z)} u(z) \quad (\text{A.11})$$

$$= \frac{b_{21}z^{-1} + (a_{21}b_{11} - a_{11}b_{21})z^{-2}}{1 - (a_{11} + a_{22})z^{-1} + (a_{11}a_{22} - a_{12}a_{21})z^{-2}}. \quad (\text{A.12})$$

Appendix B

Mini-Mast state-space model

The continuous-time state-space model of the Mini-Mast is given by

$$\dot{\mathbf{x}}(t) = \begin{pmatrix} \mathbf{A}_1 & \mathbf{O} \\ & \ddots \\ \mathbf{O} & \mathbf{A}_5 \end{pmatrix} \mathbf{x}(t) + \begin{pmatrix} \mathbf{B}_1 \\ \vdots \\ \mathbf{B}_5 \end{pmatrix} \mathbf{u}(t) \quad (\text{B.1})$$

where

$$\mathbf{A}_1 = \begin{pmatrix} 0 & 1 \\ -28.920733 & -0.347406 \end{pmatrix}, \quad (\text{B.2})$$

$$\mathbf{A}_2 = \begin{pmatrix} 0 & 1 \\ -28.839048 & -0.228771 \end{pmatrix}, \quad (\text{B.3})$$

$$\mathbf{A}_3 = \begin{pmatrix} 0 & 1 \\ -729.718377 & -3.873707 \end{pmatrix}, \quad (\text{B.4})$$

$$\mathbf{A}_4 = \begin{pmatrix} 0 & 1 \\ -1477.941136 & -1.829934 \end{pmatrix}, \quad (\text{B.5})$$

$$\mathbf{A}_5 = \begin{pmatrix} 0 & 1 \\ -1501.392005 & -0.774956 \end{pmatrix}, \quad (\text{B.6})$$

$$\mathbf{B}_1 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & -0.006166 & 0 \end{pmatrix}, \quad (\text{B.7})$$

$$\mathbf{B}_2 = \begin{pmatrix} 0 & 0 & 0 \\ -0.004122 & 0 & 0 \end{pmatrix}, \quad (\text{B.8})$$

$$\mathbf{B}_3 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0.194500 \end{pmatrix}, \quad (\text{B.9})$$

$$\mathbf{B}_4 = \begin{pmatrix} 0 & 0 & 0 \\ -0.002723 & -0.002723 & 0 \end{pmatrix}, \quad (\text{B.10})$$

$$\mathbf{B}_5 = \begin{pmatrix} 0 & 0 & 0 \\ 0.002549 & -0.002549 & 0 \end{pmatrix}. \quad (\text{B.11})$$

The output matrix for the set of displacement sensors at the tip of the mast is

$$\mathbf{C} = \begin{pmatrix} 4.846400 & 0 & -5.821079 & 0 & 4.846400 & 0 & 0.544624 & 0 & 1.069679 & 0 \\ -0.798394 & 0 & 5.784700 & 0 & 4.911925 & 0 & -1.740127 & 0 & -1.302644 & 0 \\ -3.724298 & 0 & -0.288348 & 0 & 4.633496 & 0 & -1.597996 & 0 & -0.142804 & 0 \end{pmatrix}. \quad (\text{B.12})$$

The \mathbf{D} matrix is

$$\mathbf{D} = \mathbf{0}. \quad (\text{B.13})$$

Bibliography

- [1] Michèle Basseville. Detecting changes in signals and systems—a survey. *Automatica*, 24, No. 3:309–326, 1988.
- [2] Richard V. Beard. *Failure Accommodation in Linear Systems Through Self-Reorganization*. PhD thesis, Department of Aeronautics and Astronautics, MIT, Feb. 1971.
- [3] Chi-Tsong Chen. *Linear System Theory and Design*. Holt, Rinehart and Winston, Inc., 1984.
- [4] Edward Y. Chow. *A Failure Detection System Design Methodology*. PhD thesis, Department of Electrical Engineering, MIT, Feb. 1981.
- [5] Jean R. Dutilloy. *Generalized Parity Relations for Large Space Structures with Uncertain Parameters*. Master's thesis, Department of Aeronautics and Astronautics, MIT, Jan. 1986.
- [6] Harold L. Jones. *Failure Detection in Linear Systems*. PhD thesis, Department of Aeronautics and Astronautics, MIT, Aug. 1973.
- [7] Lennart Ljung. *System Identification: Theory for the User*. Prentice-Hall, Inc. Englewood Cliffs, 1987.

- [8] Xi-Cheng Lou, Alan S. Willsky, and George C. Verghese. Optimally robust redundancy relations for failure detection in uncertain systems. *Automatica*, 22(3):333–344, 1986.
 - [9] Lennart Ljung. *System Identification Toolbox*. The MathWorks, Inc., 1988.
 - [10] M.-A. Massoumnia. *A Geometric Approach to Failure Detection and Identification in Linear Systems*. PhD thesis, Department of Aeronautics and Astronautics, MIT, Feb. 1986.
 - [11] M.-A. Massoumnia and Wallace E. Vander Velde. Generating parity relations for detecting and identifying control system component failures. *AIAA Journal of Guidance, Control and Dynamics*, 11, no.1:60–65, Jan.-Feb. 1988.
 - [12] Walter C. Merrill. Sensor failure detection for jet engines using analytical redundancy. *AIAA Journal of Guidance, Control and Dynamics*, 8, No.6:673–682, Nov.-Dec. 1985.
 - [13] Richard Pappa et al. *Mini-Mast CSI Testbed User's Guide*. Mar. 1989.
 - [14] Alan S. Willsky. A survey of design methods for failure detection. *Automatica*, 12:601–611, Nov. 1976.
 - [15] Marc Moonen et al. On- and off-line identification of linear state-space models. *International Journal of Control*, 49, No.1:219–232, 1989.
 - [16] John E. White and Jason L. Speyer. Detection filter design: spectral theory and algorithms. *IEEE Transactions on Automatic Control*, AC-32(7):593–603, July 1987.
-

